RuggedRouter™

RX1000/RX1100 User Guide

Rugged Operating System™
on Linux

RuggedCom Inc.
30 Whitmore Road,
Woodbridge, Ontario, Canada
L4L 7Z4

Web: www.ruggedcom.com
Tel: (905) 856-5288
Fax: (905) 856-1995
Toll Free: (888) 264-0006

RUGGEDROUTER™ USER GUIDE

FOR USE WITH RX1000/RX1100 PRODUCTS

Version 1.12.6 – May 14th , 2008

RuggedCom

30 Whitmore Road

Woodbridge, Ontario

Canada L4L7Z4

Tel: (905) 856-5288

Fax: (905) 856-1995

Toll Free: (888) 264-0006

support@ruggedcom.com

http://www.ruggedcom.com

Disclaimer

RuggedCom Inc. makes no warranty of any kind with regard to this material.

RuggedCom shall not be liable for errors contained herein or for consequential damages in connection with the furnishing, performance, or use of this material.

Warranty

Five (5) years from date of purchase, return to factory. For warranty details, visit www.ruggedcom.com or contact your customer service representative.

About this User Guide

This guide is concerned with aiding the user in the configuration and operation of the RuggedRouter™ using the RuggedCom command line, setup menu and web management interfaces. Specifically, this guide details aspects of:

- Accessing the User Interfaces
- Security
- Configuring the router
- Status determination
- Performance measurement
- Uploading and downloading files
- Dealing with alarms

This guide also details operation of the RX1100 Gauntlet security appliance.

This guide is intended solely for the purpose of familiarizing the reader with the ways that the RuggedRouter™ can be used to support Routing over Ethernet, T1/E1, T3 ADSL, DDS and Frame Relay as well as act as a Serial server and time synchronization device.

Applicable Firmware Revision

This guide is applicable to RuggedRouter™ ROX 1.12.6 software revision.

Who Should Use This User Guide

This guide is to be used by network technical support personnel who are familiar with the operation of networks. Others who might find the book useful are network and system planners, system programmers and line technicians.

How To Use This User Guide

Each chapter has been prepared with a feature description, an application section and a description of the default mode of operation. It is recommended that you use this guide along with the following applicable documents.

RuggedRouter™ Installation Guide

Rugged MediaConverter™ Installation Guide

RuggedCom Fiber Guide

Gauntlet Command and Control Center (CCC) User Manual,

Gauntlet Virtual Polling Controller (VPC) User Manual

Gauntlet System Installation Manual

Gauntlet System Best Practices

Document Conventions

This publication uses the following conventions:

Note: *Means reader take note. Notes contain helpful suggestions or references to materials not contained in this guide.*

Helpful Hint
This type of note often indicates useful shortcuts or methods employed by other RuggedCom customers.

Quick Start Recommendations

The following description is included to aid those users experienced with communications equipment that may wish to attempt to configure the router without fully reading the guide.

1. Locate/mount the chassis in its final resting place and apply power.
2. The router can be configured through its web management interface, or for advanced users, through ssh. The default Ethernet addresses for ports one through four are 192.168.1.1 through 192.168.4.1. Two shell accounts, rrsetup and root, are provided. Both accounts have a default password of "admin". The web management interface uses the root account password. The rrsetup account provides a shell that configures such items as passwords, addresses, date/time and services offered by the router. The root account provides a full shell.
3. Attach a PC running terminal emulation software to the RS232 port and apply power to the chassis (default baud rate, data bits, parity - "38400 8 n 1", no hardware/software flow control). Set the terminal type to VT100. Press ENTER to obtain a login prompt.

Initial Configuration Before Attaching To The Network

4. Login as the rrsetup user with password "admin".
5. **Change the root and rrsetup passwords from the shell. Record the passwords in a secure manner.** If Radius authentication will be employed, configure at least one authentication server address.
6. Configure the router's hostname, IP address, subnet mask, and gateway addresses for the built-in Ethernet ports.
7. For an RX1100 router, the Gauntlet Security application may be configured with the passphrase allocated to the network the network address of the Command and Control Center (CCC). Note that you must also configure and activate the firewall before using the Gauntlet.
8. Ensure that the date, time and timezone fields are correctly set.
9. If Web or SSH services will not be used, these can be disabled from the setup shell.

10. All further configuration is accomplished through the web management interface. Attach the configuring host to one of the Ethernet ports configured above. Point your web browser at the address for that port, use https and specify a port number of 10000, e.g. https://192.168.1.1:10000 (or otherwise if configured in step 4). Login with the root user and password (configured above). If radius authentication is configured and a server is available, you may also login via a radius user.

Basic Web Based Configuration

11. Change the router password from the **System** menu, **Change Password** sub-menu.
12. If you are using the web management interface you may wish to restrict the allowed users to a specific subnet. This can be done in the **Webmin** menu, **Webmin Configuration**, **IP Access Control** sub-menu.
13. If you are planning to SSH in to the router you may wish to restrict the allowed users to a specific subnet. This can be done in the **Servers** menu, **SSH Server**, **Networking** sub-menu.
14. The router's local hostname may configured in the **System Menu**, **System Hostname** sub-menu.
15. The router may be configured to log to a remote server by the **Maintenance menu**, **System Logs** sub-menu. See the chapter "Maintaining The Router" for more details.
16. The router's DNS settings may configured in the **DNS Clients** sub-menu. You may also specify the IP addresses of frequently used hosts. See the chapter "Configuring Networking" for more details.

Physical Interface Related

17. Ethernet port parameters may be changed in the **Networking** menu, **Ethernet** sub-menu. The **Ethernet Interfaces** sub-menu will configure the IP address, subnet mask, gateway address, proxy arping and media type of each interface. See the chapter "Configuring Ethernet Interfaces" for more details.
18. If your router is equipped with T1/E1 WAN interfaces, the **Networking** menu, **T1/E1** sub-menu will allow you to configure them with Frame Relay or PPP connections. See the chapter "Configuring Frame Relay/PPP And T1/E1" for more details.
19. If your router is equipped with T3 WAN interfaces, the **Networking** menu, **T3** sub-menu will allow you to configure them with Frame Relay or PPP connections. See the chapter "Configuring Frame Relay/PPP And T3" for more details.
20. If your router is equipped with DDS interfaces, the **Networking** menu, **DDS** sub-menu will allow you to configure them with Frame Relay or PPP connections. See the chapter "Configuring Frame Relay/PPP And DDS" for more details.
21. If your router is equipped with ADSL interfaces, the **Networking** menu, **ADSL** sub-menu will allow you to configure them. See the chapter "Configuring PPPoE On ADSL" for more details. If you wish to use PPPOE with an external ADSL modem, the **Networking** menu, **Ethernet** sub-menu will configure it.

22. If your router is equipped with an embedded modem, the **Networking** menu, **Modem** sub-menu will allow you to configure it with PPP or incoming console connections. See the chapter "Configuring PPP And Modem" for more details.
23. If your router is equipped with Serial Interfaces, the **Servers** menu, **Serial Protocols** sub-menu will allow you to configure them with an operating protocol. See the chapter "Configuring Serial Protocols" for more details.
24. If your router is equipped with a Precision Time Protocol Card, the **Servers** menu, **IRIGB** sub-menu will allow you to enable and configure its output ports. See the chapter "Configuring IRIGB" for more details.

Additional Configuration

25. You may wish to configure a backup interface to use in the event of a failure of your default gateway interface. This can be done in the **Networking** menu, **Network Configuration, End To End Backup** sub-menu.
26. If you are planning to connect your router to the Internet, configure the firewall and then activate it. This can be done in the **Networking** menu, **Shorewall Firewall** sub-menu.
27. The router provides a default event logging configuration. You can modify this configuration through the **Maintenance** menu, **System Logs** sub-menu. Remote logging can be activated here.
28. The routers SSH and Web Management interfaces are enabled by default. The routers DHCP server, IPsec VPN server, NTP server, OSPF/RIP protocol, VRRP protocol and firewall are disabled by default. To changes these services visit the **System** menu, **Bootup and Shutdown** sub-menu.
29. You can install static IP and Multicast routings for Ethernet and WAN interfaces via the **Networking** menu, **Network Configuration**, **Routing and Default Route** and **Static Multicast Routing** sub-menus.
30. You can configure the NTP server through the **Servers** menu, **NTP Server** sub-menu. See the chapter "Configuring NTP" for more details.
31. You can configure SSH through the **Servers** menu, **SSH Server** sub-menu. SSH can be set-up to issue a login banner from this menu. See the chapter "Configuring SSH" for more details.
32. Traffic prioritization can be configured on the network interfaces through the **Networking** menu, **Traffic Prioritization** sub-menu.. See the chapter "Traffic Prioritization" for more details.
33. SNMP is disabled by default. You can configure SNMP by following the instructions in the Appendix on SNMP. You may allow read and write access, set community names, enable traps and program the router to issue traps with a specific client address.
34. If your router is an RX1100 you may configure and activate the Snort Intrusion Detection system and the Gauntlet Security Appliance. If you decide to forward daily email summaries you must configure a mail forwarder in the **Maintenance** menu **Miscellaneous** sub-menu **Outgoing Mail** sub-menu.

35. When your routers configuration is stable, it is recommended that the configuration should be uploaded from the router and stored as a backup. The **Maintenance** menu **Backup And Restore** sub-menu will be useful.
36. Should you need to transfer files to or from the router, the **Maintenance** menu **Upload/Download Files** sub-menu will be useful.
37. Further concerns such as ensuring robustness, measuring and optimizing performance are dealt with by reading the guide fully.

Table Of Contents

Table Of Figures

Chapter 1 – Setting Up And Administering The Router

Introduction

This chapter familiarizes the user with the RuggedCom Serial Console interface, the RuggedRouter™ Setup script and signing on to the Web interface. This chapter describes the following procedures:

- Running the Setup Script
- Signing on the Web Interface
- Signing on to the Command Prompt
- Restoring the default configuration

Access Methods

You can access the router through the console, Ethernet ports, WAN ports and the modem port.

Accounts And Password Management

The router provides an "rrsetup" account which provides a shell that quickly configures such items as passwords, addresses, date/time and services offered by the router. It is very useful to sign-in to this shell first, harden the router, and configure network addresses in order that the router be reachable from the network through Web Management. **The rrsetup password should be changed, recorded securely and restricted to qualified personnel.**

The root account provides a superuser capability for SSH shell access and the Web server. **The password should be changed, recorded securely and restricted to qualified personnel.**

The root and rrsetup accounts may be also be managed through radius authentication.

The Web management agent can be accessed through the root account. It may also be accessed through a number of radius accounts via radius authentication. This offers the advantage of attributing actions in logs to the specific user, as opposed to the root user.

Default Configuration

Your RuggedRouter™ is shipped from the factory with the following defaults:

- Ethernet ports are enabled and have an address of 192.168.X.1 where X is the port number,
- WAN and modem ports are disabled,
- IRIG-B output ports are disabled,
- Setup account "rrsetup", password "admin",
- Superuser account "root", password "admin",
- SSH and Web Management interfaces are enabled by default. All other services (including Serial Protocol Server, DHCP server, NTP server, End to

End Backup Server, VPN Server, NFS, OSPF/RIP protocol and firewall) are disabled by default.

Accessing The RuggedRouter™ Command Prompt

From the Console Port

Attach a terminal (or PC running terminal emulation software) to the RS232 port on the rear of the chassis. The terminal should be configured for 8 bits, no parity operation at 38.4 Kbps. Hardware and software flow control must be disabled. Select a terminal type of VT100.

Once the terminal is connected, pressing <CR> will prompt for the user to login as and that user's password. Sign-in as either the rrsetup or root user. The router is shipped with default passwords of "**admin**" for either of these accounts.

From SSH

Use an SSH agent running the version 2 protocol. SSH to either the rrsetup or root accounts of the router at one of its IP addresses described above. The router is shipped with default passwords of "**admin**" for either of these accounts.

The RuggedRouter Setup Shell

Signing-in as the rrsetup user will automatically enter the configuration shell shown below. Quitting the shell (with cancel, or by entering escape) will cause the connection to close.

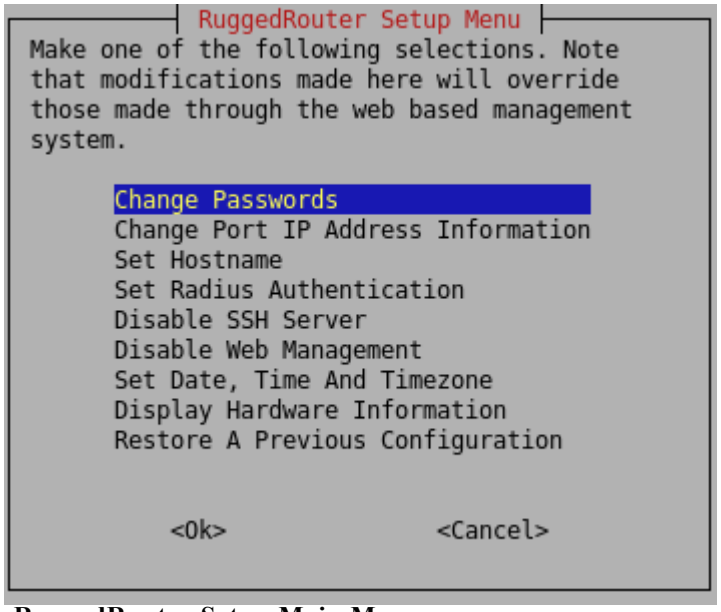

Figure 1: RuggedRouter Setup Main Menu

The shell provides a number of configuration commands, described below.

Configuring Passwords

The **Change Passwords** command changes the rrsetup and root account passwords. These passwords should be changed before installing the router on the network.

Figure 2: RuggedRouter Setup Password Change Menu

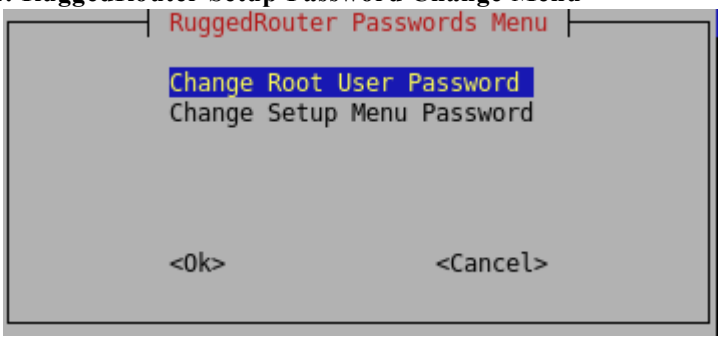

Configuring IP Address Information

The **Change Port IP Address** command configures port IP addresses and gateways.

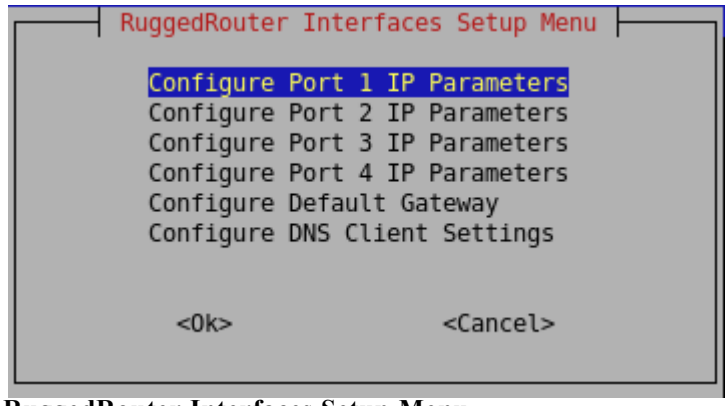

Figure 3: RuggedRouter Interfaces Setup Menu

Each port number X has a default address of 192.168.X.1 and a mask of 255.255.255.0.

The **Configure Default Gateway Settings** command configures the default gateway.

The **Configure DNS Client Settings** command configures the DNS server address. If the router is part of a domain, enter the domain name in the "Search Domain" field.

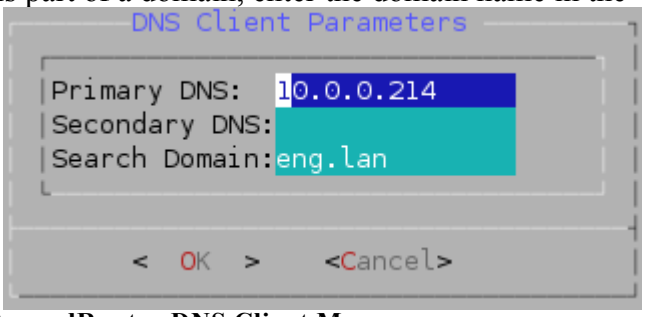

Figure 4: RuggedRouter DNS Client Menu

Setting The Hostname

The **Set Hostname** command sets the hostname, shown in shell prompts and Web Management.

Configuring Radius Authentication

The **Set Radius Authentication** command configures the address of a Radius server, if available.

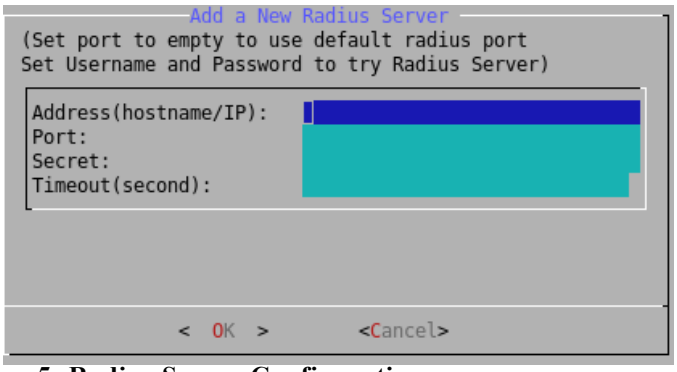

Figure 5: Radius Server Configuration menu

The **Hostname/IP** and **Port Number** fields configures the server location.

The **Shared Secret** field configures the unique password used by this server.

The time **Timeout** field selects the maximal time to wait before trying the next server.

The entry, created for both LOGIN and PPP Login, can be changed from the web interface.

Enabling And Disabling The SSH and Web Server

By default SSH and Web Management are enabled. The **Disable SSH** and **Disable Web Management** commands allows these services to be disabled. The servers will be immediately stopped. If access to the shell has been made through ssh the session will continue, but no new sessions will be allowed.

Upon disabling the services, the titles in the main menu will change to **Enable SSH** and **Enable Web Management** to reflect the disabled state. Enabling a service automatically restarts it.

Enabling And Disabling The Gauntlet Security Appliance

The Gauntlet security Appliance requires a pass phrase unique to your network. This menu will configure it.

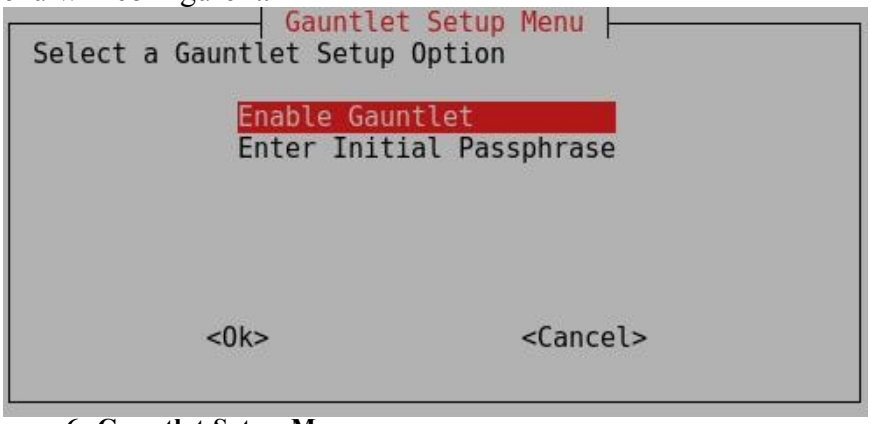

Figure 6: Gauntlet Setup Menu

Configuring The Date, Time And Timezone

The **Set The Date, Time And Timezone** command allows these parameters to be set.

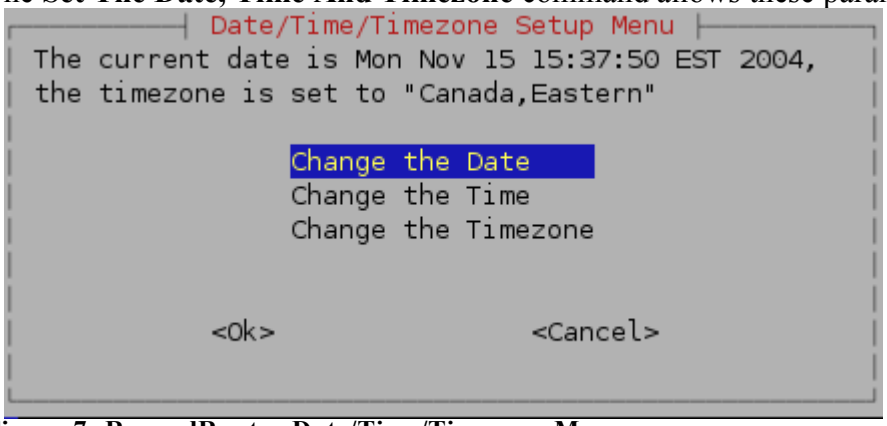

Figure 7: RuggedRouter Date/Time/Timezone Menu

Once set, the router will account for Daylight Savings time.

Displaying Hardware Information

The Display Hardware Information command describes commissioned hardware.

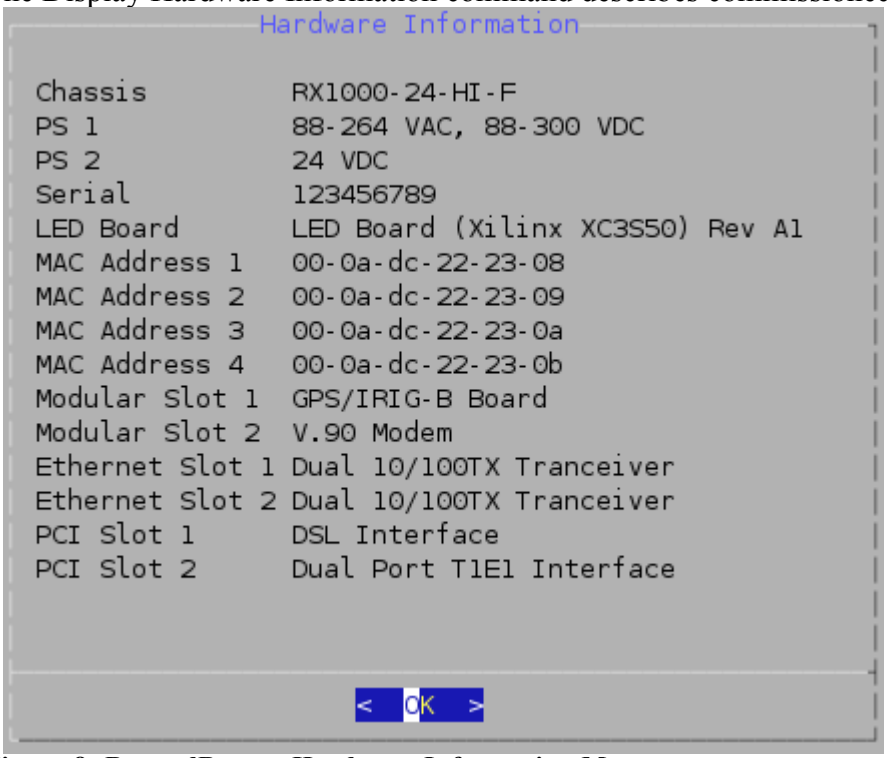

Figure 8: RuggedRouter Hardware Information Menu

Restoring A Configuration

The **Restore A Previous Configuration** command provides a means to restore a previously taken snapshot of the configuration of the router.

Note: The router will reboot immediately after restoring configuration.

The user is first prompted to select either the factory default configuration or a previously made archive.

Note: Restoring the factory defaults will reset IP addresses and may make the router impossible to reach from the network.

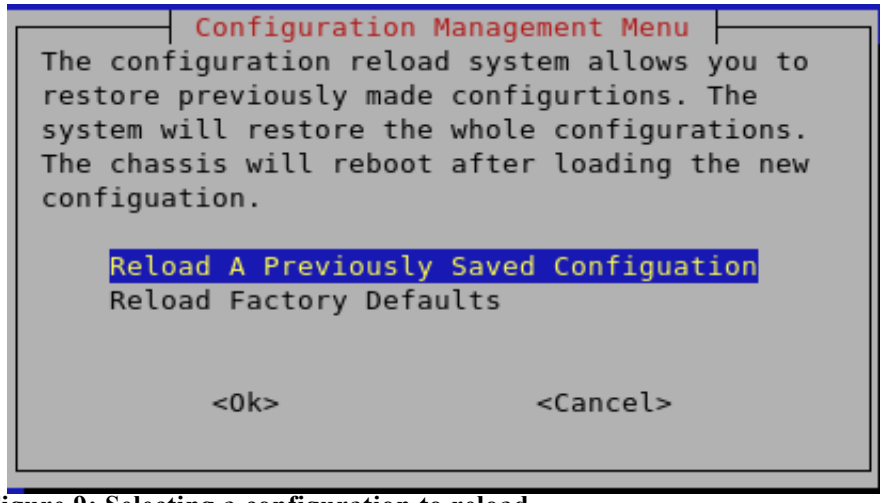

Figure 9: Selecting a configuration to reload

Initially, your RuggedRouter will have no previously saved configurations. The factory defaults will always be available.

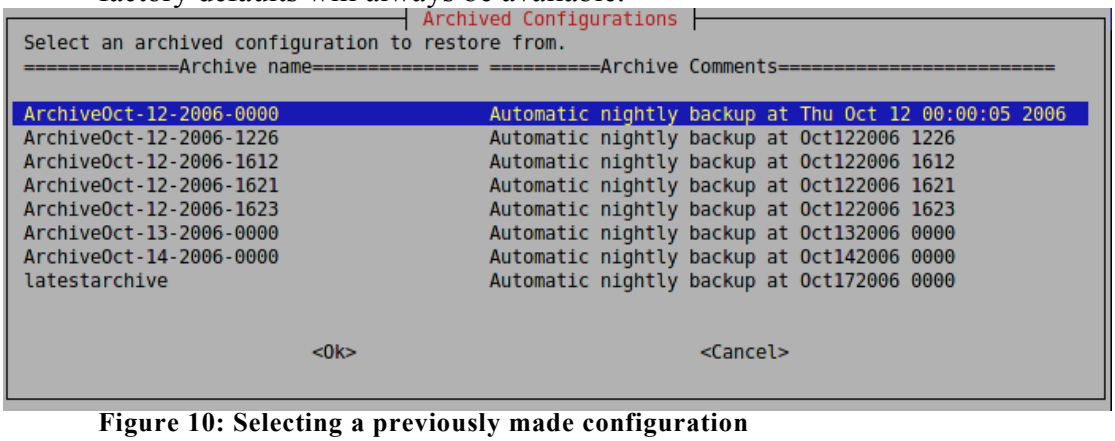

Figure 10: Selecting a previously made configuration

Once a configuration is selected the archive will be restored. After the configuration is restored, the router will reboot immediately.

The RuggedRouter™ Web Interface

The RuggedCom Web interface is provided by an enhanced version of the popular Webmin interface.

Using a Web Browser to Access the Web Interface

Start a web browser session and open a connection to the router by entering a URL that specifies its hostname or IP address (e.g. https://179.1.0.45:10000). Once the router is contacted, start the login process by clicking on the "Login" link. The resulting page should be similar to that presented below.

Figure 11: Signing On To The Router With A Web Browser

Enter the "root" user name and the appropriate password for that user, then click on the "Login" button. The router is shipped with a default administrator password of "admin". Once successfully logged in, the user will be presented with the main menu.

SSL Certificate Warnings

Your browser may complain about the SSL certificate that Webmin issues.

This happens because the default SSL certificate that comes with Webmin is not issued by a recognized certificate authority. From a security point of view, this makes the certificate less secure because an attacker could theoretically redirect traffic from your server to another machine without you knowing, which is normally impossible if using a proper SSL certificate.

Network traffic is still encrypted though, so you are safe against attackers who are just listening in on your network connection.

If you are initiating the connection to the router, and your network is private, a VPN or firewalled, it should be safe to have your browser permanently accept the certificate.

If you want to be really sure that the Webmin server you are connecting to is really your own, the only solution is to order a certificate from an authority like Verisign that is associated with your router's hostname and will be recognized by web browsers.

The Structure of the Web Interface

The Web interface presents an web page with two frames. The leftmost or index frame selects subsystems to configure and is always displayed.

The rightmost or configuration frame presents the configuration for the currently selected subsystem, or in the case of signing-on, the home page window. The home page window presents an annotated view of the front of the chassis as well as a number of important system parameters. These parameters include:

- The router uptime and load averages for the past 1, 5 and 15 minutes. Under normal operation the load average should be less than 2.0.
- The disk usage. A disk usage higher than 92% requires attention.
- The memory usage, indicating the amount of memory used by applications. Under normal operation memory usage should be less than 60%.
- The chassis temperature.
- Any major alarms, such as the failure of hardware components.

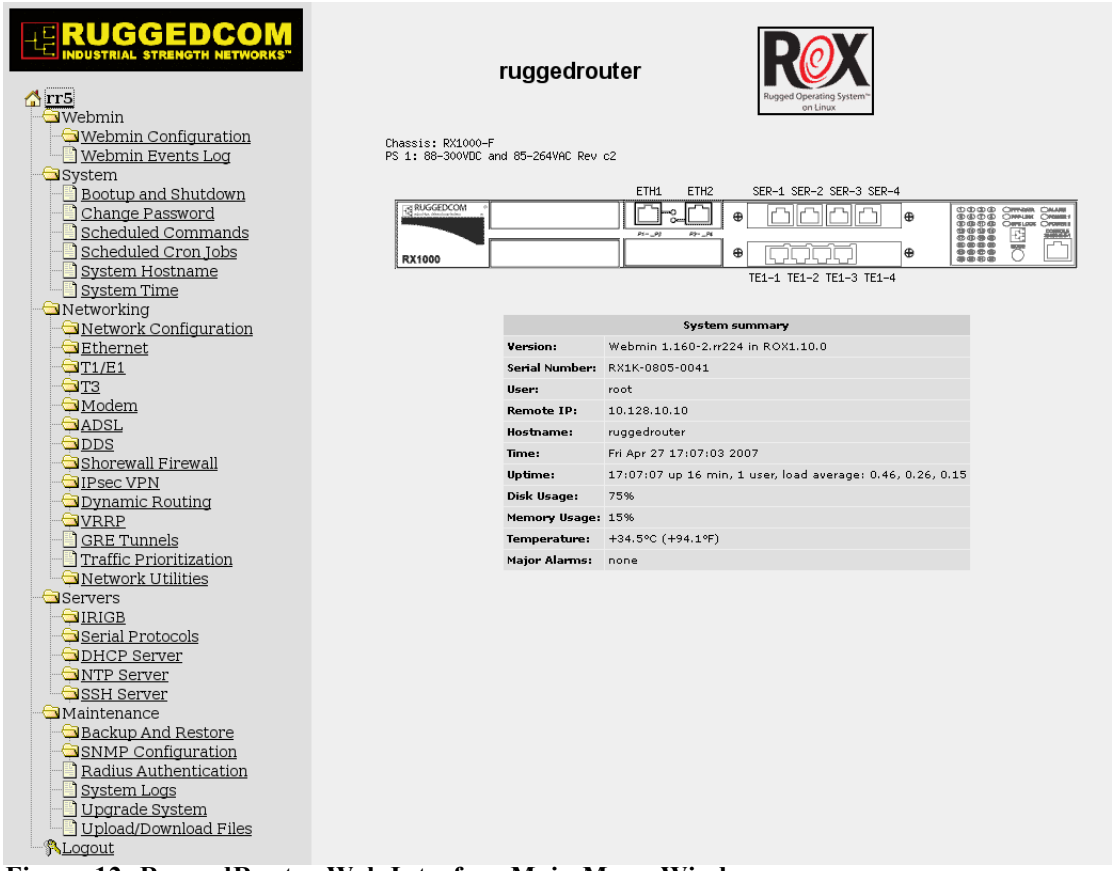

Figure 12: RuggedRouter Web Interface Main Menu Window

The index frame presents a number of entries with associated icons:

- The icon forces home page window to be redisplayed.
- The icon signifies that the next level contains a menu of menus.
- The icon signifies that clicking the entry will run a single menu.
- The icon logs out of Webmin.

The menu system entries are composed of the Webmin, System, Servers, Networking and Maintenance menus.

The Webmin Menu provides the ability to:

- Configure the sign-on password,
- Specify session timeouts,
- Restrict the Subnet of IP addresses that can login,
- Configure and view Webmin event logs,

The System Menu provides the ability to:

- Change the router password,
- Enable and disable applications from running,
- Reboot the router,
- Schedule one time and periodic tasks to run,
- Change the router's name (hostname),
- Change the time and date.

The Servers Menu provides the ability to:

- Control and configure the Serial Protocol, DHCP, NTP, IRIGB and SSH servers.

The Networking Menu provides the ability to:

- Configure the network interfaces,
- Configure static IP and Multicast Routings and configure a default gateway,
- Select a DNS server and edit local host addresses,
- Configure End To End Backup,
- Configure DDS, T1/E1, T3 and ADSL Networking,
- Configure the embedded modem,
- Set up the firewall,
- Set up Virtual Private Networking,
- Configure Routing protocols such as OSPF and RIP,
- Configure Virtual Router Redundancy Protocol (VRRP),
- Configure Traffic Prioritization,
- Perform pings, traceroutes, host lookups and line tracing.

The Maintenance Menu provides the ability to:

- Manage the Gauntlet Security Appliance
- Backup and restore configurations,
- Configure SNMP access,
- Configure Radius Authentication,
- View system logs,
- Upgrade the software of the router,
- Upgrade the router type to RX1100,
- Upload/Download files to and from the router.

Using The LED Status Panel

Figure 13: LED Status Panel

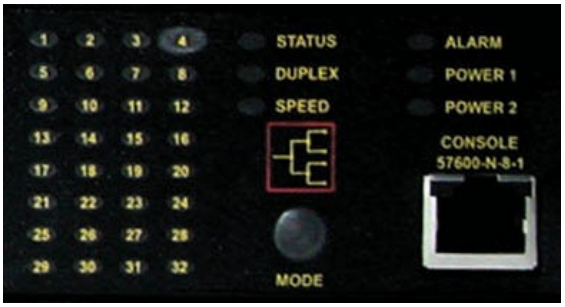

The LED status Panel provides the console port, indicates the status of hardware/software and can initiate a controlled reboot.

The LEDs are organized into three primary groups; the port group, GPS/PPP group and the Alarm/Power Supply group. The display possibilities are as follows:

LED Name	Description
LED 1-4	Ethernet port 1-4 is active when green
LED 5-8	Ethernet port 1-4 has link when green and failed when red
LED 9-12	WAN port 1-4 is active when green
LED 13-16	WAN port 1-4 has link when green and failed when red
LED 17-20	WAN port 5-8 is active when green
LED 21-24	WAN port 5-8 has link when green and failed when red
PPP-DATA	PPP Modem port is active when green
PPP-LINK	PPP Modem port has link when green
GPS-LOCK	The PTP card GPS system has satellite lock
ALARM	A Major Alarm exists when red
POWER 1	Power supply 1 working properly when green and failed when red
POWER 2	Power supply 2 working properly when green and failed when red

Figure 14: Meaning of LEDs

The software will cause the ALARM LED to become active for various reasons. Any condition that causes the ALARM LED to become active will activate the critical fail relay. The Web interface displays the alarms.

Pressing the pushbutton for more than five seconds will reboot the router.

Obtaining Chassis Information

The chassis displays the hardware inventory at boot time. This information is captured in the /var/log/messages file after boot. The Web Management interface home page displays the chassis serial number.

Chapter 2 – Webmin Configuration

Introduction

This chapter familiarizes the user with configuring the router through the Webmin menu and describes the following procedures:

- Configuring the IP Address and Subnet Mask
- Configuring the Gateway Address
- Viewing the Webmin Log

Webmin Configuration Menu

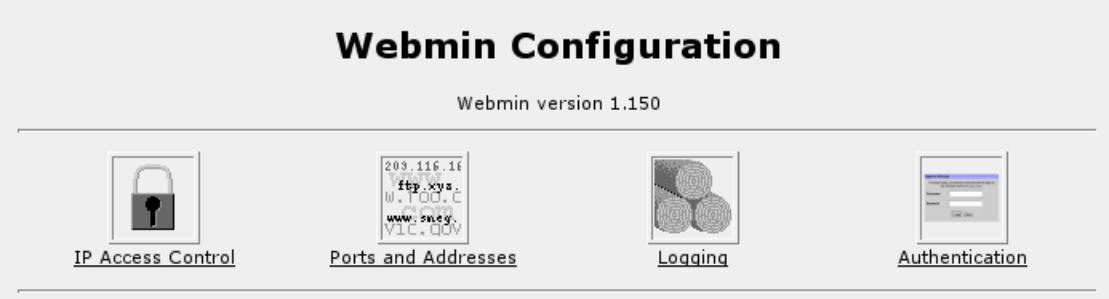

Figure 15: Webmin Configuration Menu

IP Access Control

Figure 16: Webmin Configuration Menu, IP Access Control

Webmin uses a secure communications method called Secure Sockets Layer (SSL) to encrypt traffic with its clients. Webmin guarantees that communications with the client is kept private. But Webmin will provide access to any client that provides the correct password, rendering it vulnerable to brute force attacks. The best way of addressing this problem is to restrict access to specific IP addresses or subnets.

By default, IP access control allows all IP addresses to access Webmin.

If your router is being used on a completely private network, or IP access control is being provided by the firewall you may leave IP Access Control disabled. Select the **Allow from all addresses** field and Save.

If you wish to restrict access to a single address or subnet, select the **Only allow from listed addresses** field. Enter a single IP address or a subnetted address.

If you wish to deny access to a specific subnet, select the **Deny from listed addresses** field. Enter a single IP address or a subnetted address.

If DNS is configured you may allow and deny based upon hostname. Partially qualified domain names such as *.foo.com are acceptable.

The **Resolve hostnames on every request** field forces Webmin to perform a hostname lookup for every user access. The result of this will be that a dynamically assigned IP with a DNS entry with a Dynamic DNS registrar will be able to be checked against the IP Access Control list, just like a fixed address. This method is useful for administrators who travel or simply don't have a fixed address at their normal location.

Note: This is not efficient if you have more than a few domain names entered in the IP Access Control list, due to the high overhead of performing a name lookup for every hostname in the list on every request.

Ports And Addresses

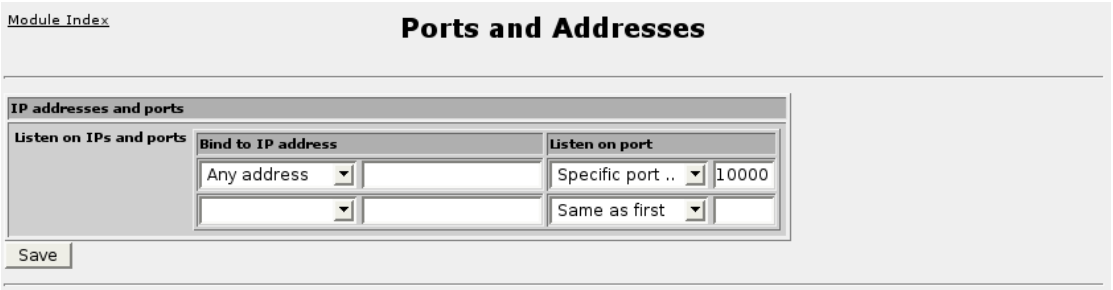

Figure 17: Webmin Configuration Menu, Ports and Addresses

This command allows you to restrict access to Webmin from one particular network interface on your server. If your Webmin server has a non-routable local address and a routable Internet address, you should decide whether anyone will ever need to be able to access the Webmin server from outside of your local network. If not, simply configure Webmin to listen on the local interface.

By default, Webmin listens on TCP port 10000 for clients. It is possible to change this default behaviour.

Change Help Server

Figure 18: Webmin Configuration Menu, Change Help Server

The Web management package provides context sensitive help in each of its menus. When a help link is selected the router instructs the browser to open the help text from a help server. In this way the router does not waste large amounts of disk space storing help text and network bandwidth sending large web pages. By default, the router directs the browser to the same server used to upgrade the router. This is as specified in the Maintenance menu Upgrade System sub-menu Change Repository Server command.

This command allows you to disable Web management help, use the upgrade repository server as well as specify a new server. If you specify an alternate web server to host the help text, you must install release specific help directories below the document root. The menu suggests the currently expected directory. The actual help files are provided with every release under the html directory at the repository server.

Logging

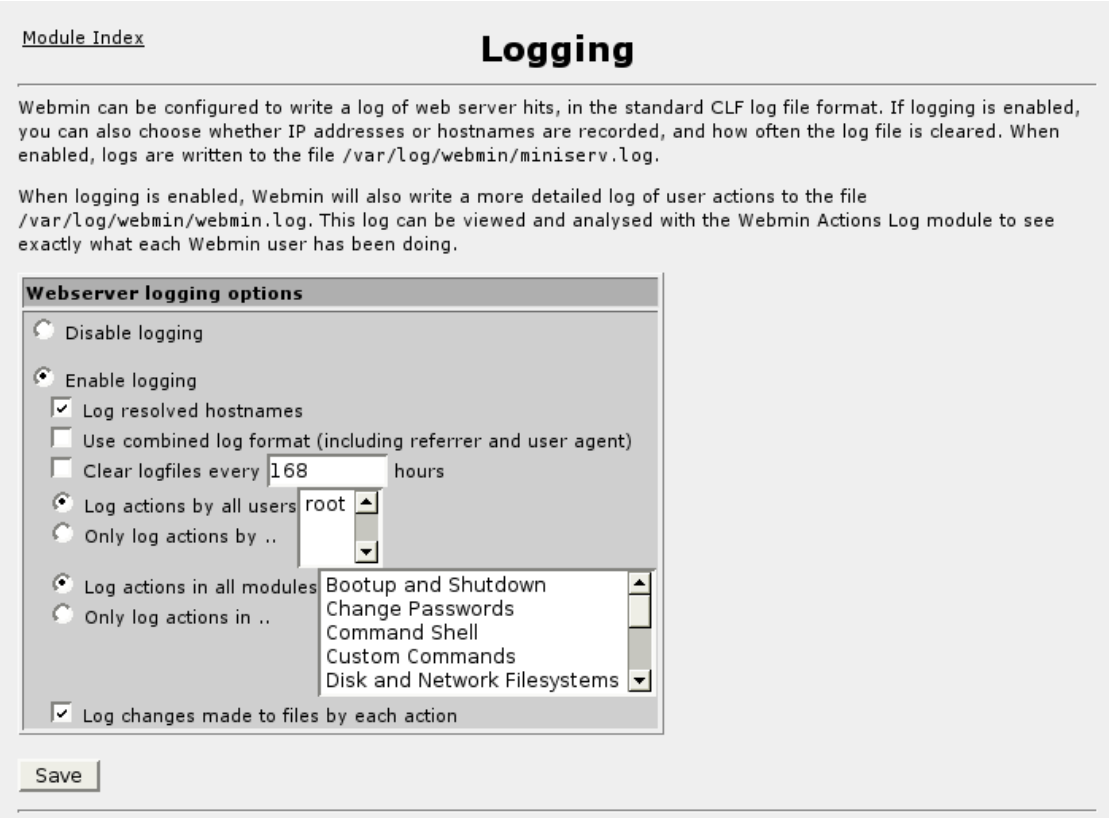

Figure 19: Webmin Configuration Menu, Logging

This menu allows you to log actions taken by Webmin administrators.

It is also possible to log actions based on the module where the actions are performed.

The **Log resolved hostnames** field will cause Webmin to provide a hostname rather than just an IP address for the client computer that performed an action.

The **Clear logfiles every...hours** field causes Webmin to rotate its own logs and keep them from overfilling the disk with old logs.

Currently, the **Log actions by all users** field should be left selected.

The **Log changes made to files by each action** field causes verbose logging and should be left enabled.

Authentication

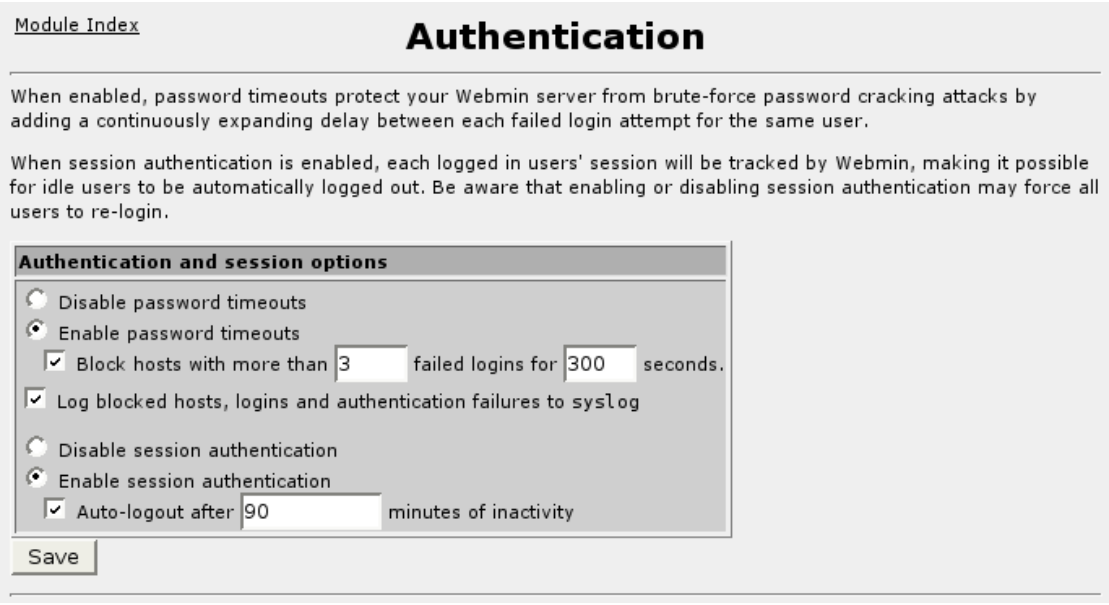

Figure 20: Webmin Configuration Menu, Authentication

This menu allows you to configure what Webmin will do when a number of failed logins from the same IP address occur.

If the **Enable password timeouts** field is selected, the host will be blocked for the specified period of time. If the **Log blocked hosts, logins and authentication failures to syslog** field is selected, warning messages will be added to the syslog.

Enabling the **Enable session authentication** field, activating "**Auto-logout after..**" will cause an individual administrators session to be logged out after the specified period.

Webmin Events Log

Figure 21: Webmin Events Log

This menu allows you to search the Webmin log for changes made by yourself or other administrators.

This page intentionally blank

Chapter 3 – Configuring The System

Introduction

This chapter familiarizes the user with:

- Enabling and disabling processes such as SSH and Web Management
- Changing The Password
- Shutting down and Rebooting the system
- Scheduling one-off and periodic commands
- Examining system logs
- Changing the hostname
- Changing the system time and timezone

Bootup And Shutdown

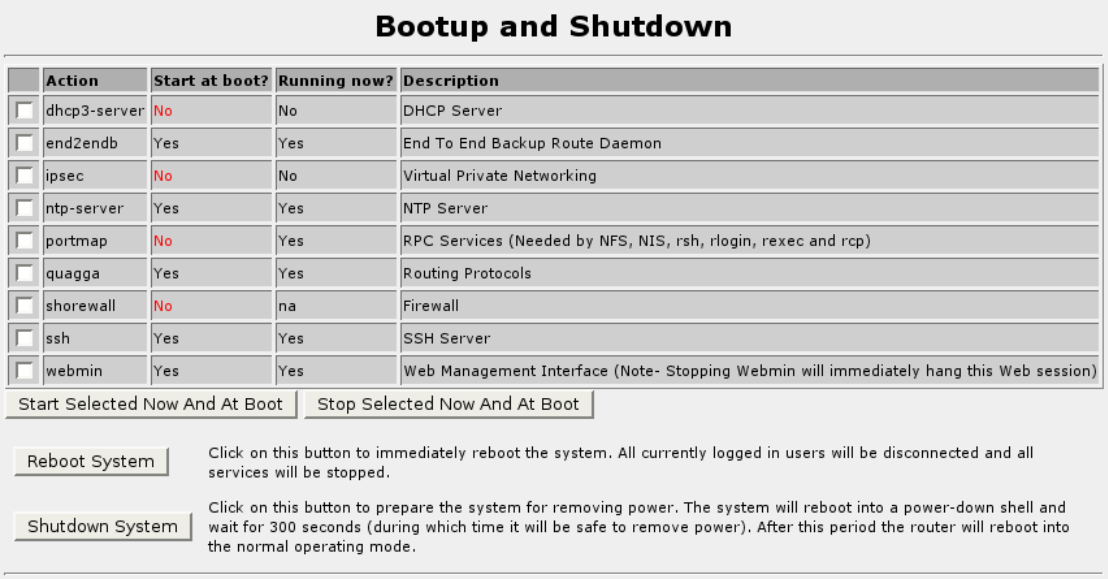

Figure 22: Bootup and Shutdown, Part 1

This menu allows you to enable/disable services and to perform actions at boot. The first part of the menu manages services. Check the box for the desired service and click on "Start Selected" to start the service and have it start at the next boot. Click on "Stop Selected" to stop the service and not have it start at boot.

The "Reboot System" button will cause the system to reboot.

The "Shutdown System" button shuts down the system in order to remove power.

Note: *The RuggedRouter* ***never*** *enters a permanent shutdown state. If the RuggedRouter is instructed to shutdown, either from Webmin or from a shell command,* ***it will reboot into a command line shell that waits five minutes before restarting****.*
If you really want the router to remain powered but permanently inactive, you must issue the shutdown, connect a terminal to the serial port, wait for the router to enter the shutdown shell and issue a CTRL-C.
Once again, if you accidentally shutdown the router it will restart after five minutes.

The second part of the menu allows you to program specific actions at boot time. The script will be run after all regular boot actions have completed.

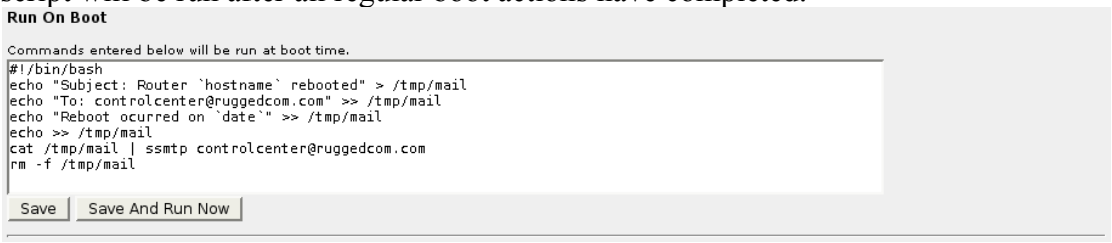

Figure 23: Bootup and Shutdown, Part 2

The actions may be a series of commands that can be executed at the command line. Each entered line is executed independently of the previous line, so change directory commands will not be effective. Always specify the absolute path of files used in commands. Selecting **Save And Run Now** will run the script and show its output, allowing you to debug it.

Change Password Command

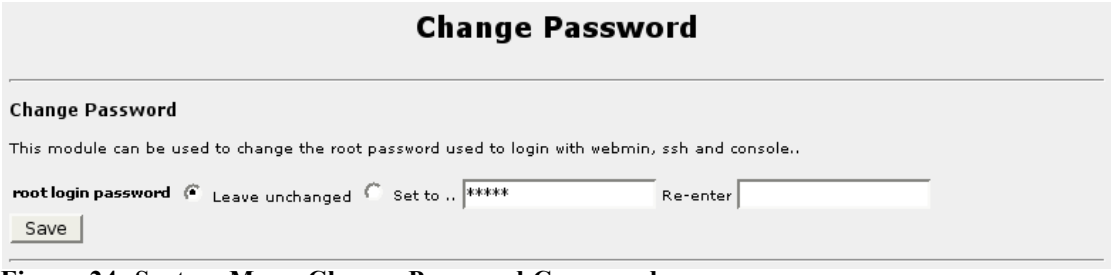

Figure 24: System Menu Change Password Command

This command changes only the root account password used to login to Webmin and the root account via the serial console or SSH.

Scheduled Commands

Figure 25: Scheduled Commands

This menu allows you to schedule a command to run in the future.

Begin by selecting the time and date you wish to run the command at using the **Run on date** and **Run at time** fields.

Use the **Run in directory** field to enter a directory to run the command in, or simply use "/".

Finally, enter the command to execute in the **Commands to execute** field.

Note that the command will remain scheduled after reboot. After the command is entered, the Scheduled Commands menu will display any commands and allow you cancel them.

Figure 26: Scheduled Commands Displaying a Command

Scheduled Cron Jobs

A Cron job is a combination of a command to run, and a definition of the times at which to run it. The Scheduled Cron Jobs allows you to create, delete and edit these jobs.

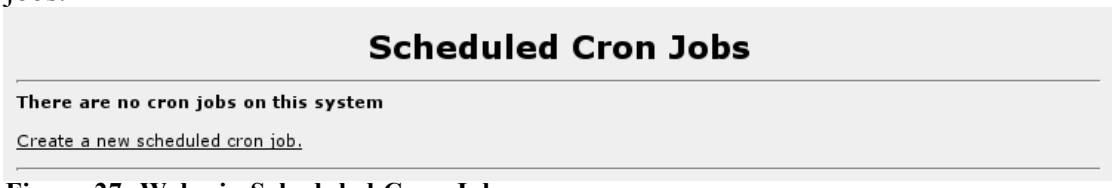

Figure 27: Webmin Scheduled Cron Jobs

Initially, there will be no scheduled jobs. Follow the "create" link to create one.

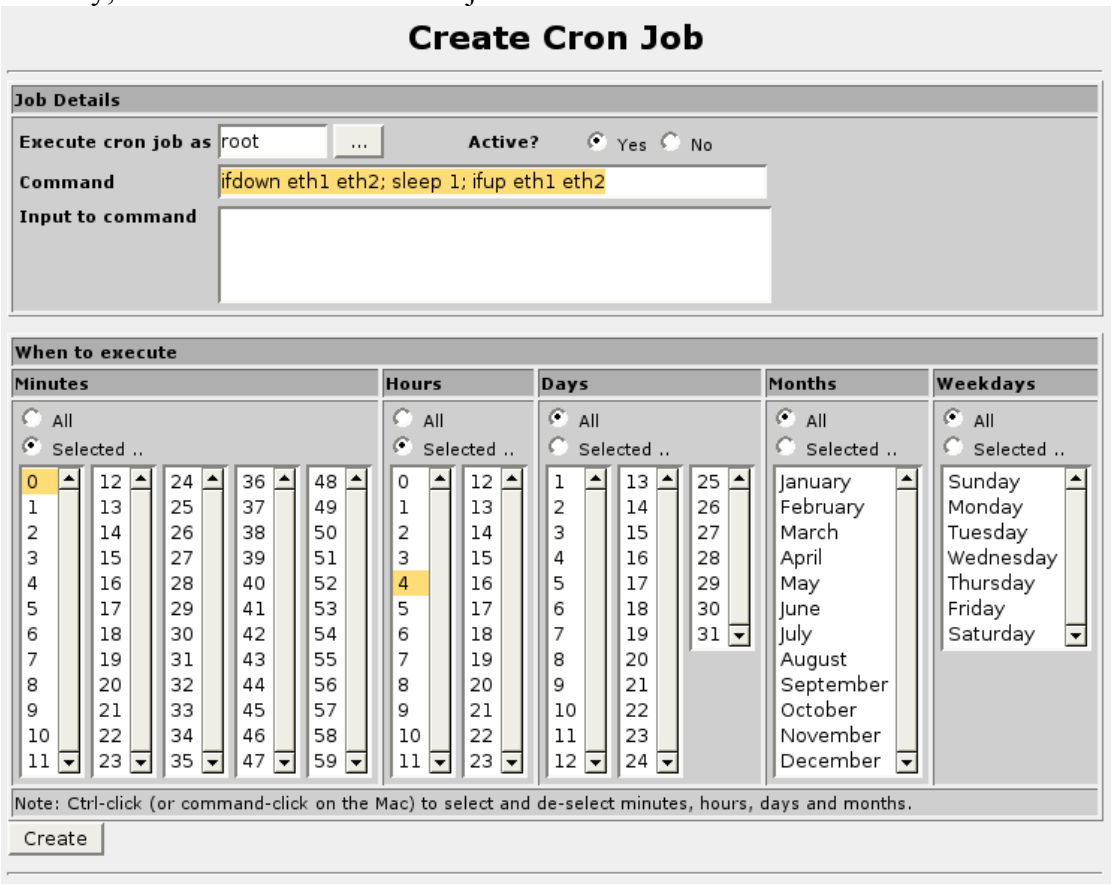

Figure 28: Creating a Cron Job

Begin the construction of the job by selecting a "user" to execute as. For most purposes, "root" will suffice. Enter this user in the **Execute cron job as** field

Enter the command to execute and any input to the command in the **Command** field. Select the times the script is to run from the **When to execute** table (remember to check the **selected** button above any column you edit).

The **Active** radio button at the top of the menu temporarily disables the job.

After selecting the **Create** button, the Scheduled Cron Jobs menu will display the job.

Figure 29: Scheduled Cron Jobs menu displaying cron jobs

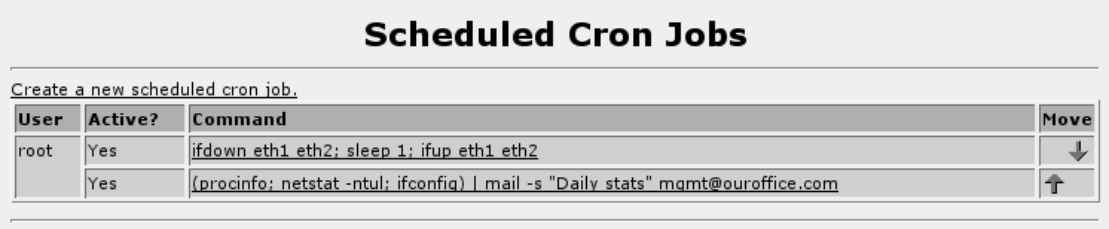

Follow the link of a specific job in order to delete the job, edit it, or test the command part of the job by running it immediately.

If you have multiple jobs, the arrows in the **Move** column will alter the order in which they are presented.

System Hostname

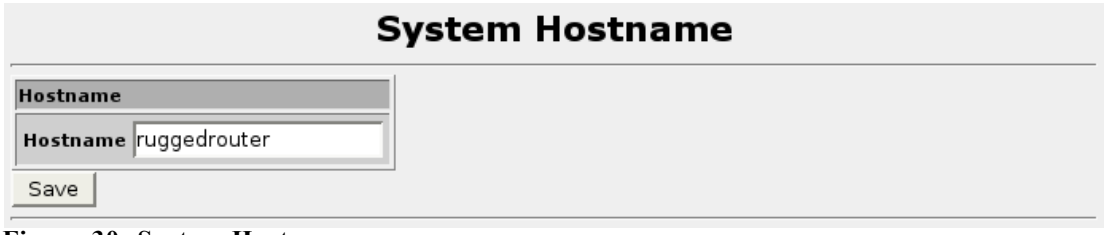

Figure 30: System Hostname

The **Hostname** field modifies the hostname as presented in the web server and shell sessions. Note that the new hostname will only appear in new sessions.

System Time

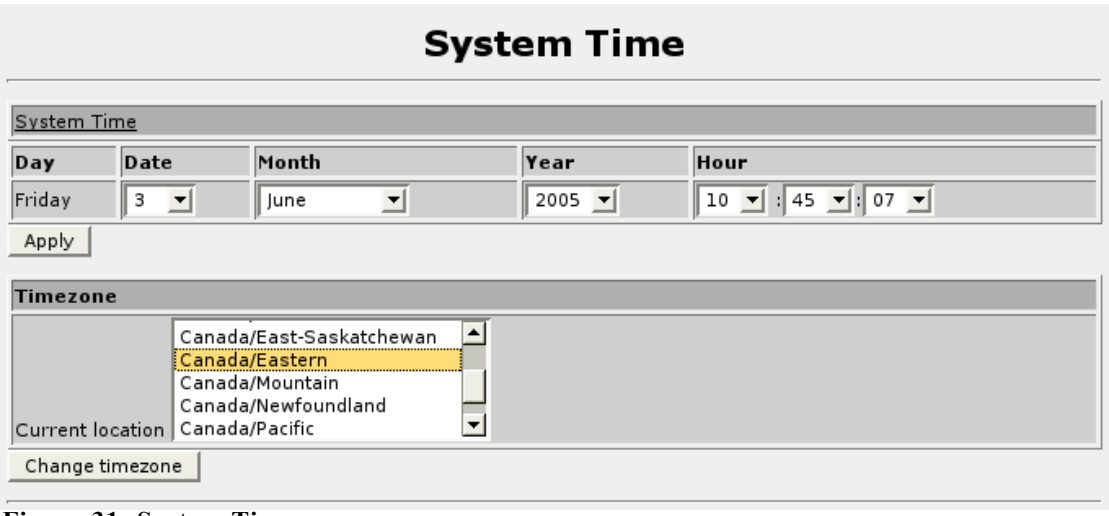

Figure 31: System Time

This menu provides a method to set the time and timezone of the router.

Note: Changing the system may confuse protocols such as OSPF and RIP, which depend upon an accurate system time. *If you use OSPF or RIP, changing the time from this menu will restart them.*

This page intentionally blank

Chapter 4 – Configuring Networking

Introduction

This chapter familiarizes the user with:

- Configuring Routing and Gateways
- Configuring DNS
- Entering host addresses
- Configuring a pair of End To End Backup Interfaces
- Viewing Routing Tables

Network Configuration

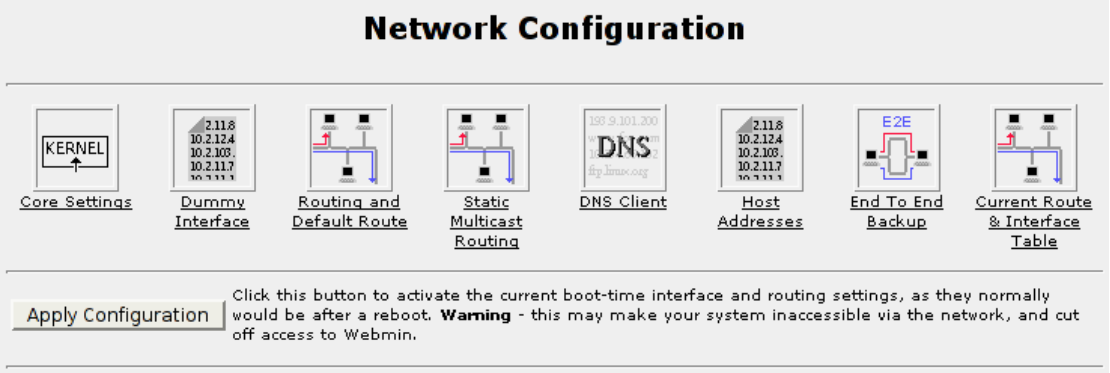

Figure 32: Network Configuration Menu

This menu allows you to configure IP networking parameters.

Select the **Core Settings** icon to configure kernel networking settings such as antispoofing and syncookies filtering.

Select the **Dummy Interface** in order to to assign an IP Address to the router that is independent of its interfaces.

Select the **Routing and Default Route** icon to assign a gateway address.

Select the **Static Multicast Routing** icon to configure static multicast routes.

Select the **DNS Client** icon to point the router at a DNS server.

Select the **Host Addresses** icon to locally configure IP address-hostname mappings.

Select the **End To End Backup** icon to configure an end to end backup connection.

Select the **Current Routing & Interface Table** icon to view the routing table.

The **Apply Configuration** button serves to restore the permanently saved changes and restart Ethernet networking.

Core Settings

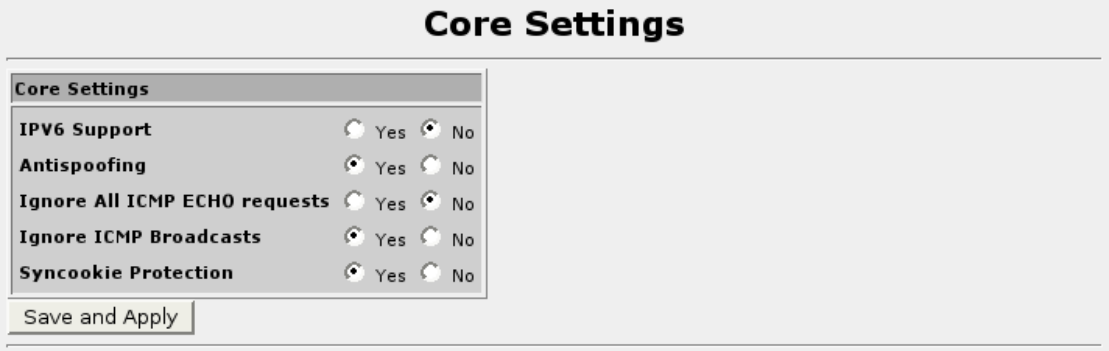

Figure 33: Core Networking Settings

This menu allows you to configure core networking settings.

The **IPV6 Support** field determines where IPV6 interfaces are created and supported at boot time. Set this option to yes if you need these interfaces. Disabling these interfaces removes them from interface displays and OSPF/RIP. A change will take effect at the next boot.

The **Antispoofing** field corresponds to the kernel rp_filter setting. Setting Antispoofing to "yes" will cause the kernel to reject incoming packets if their source address doesn't match the network interface that they're arriving on, which helps to prevent IP spoofing. If you modify this parameter, the setting be applied to all active interfaces, change the default setting for new interfaces and those created at bootup.

The **Ignore All ICMP ECHO** field corresponds to the kernel icmp_echo_ignore_all setting. Setting Ignore All ICMP ECHO to "yes" will cause the kernel to reject incoming ICMP ECHO request packets.

The **Ignore ICMP Broadcasts** field corresponds to the kernel icmp_echo_ignore_broadcasts setting. Setting Ignore ICMP Broadcasts to "yes" will cause the kernel to reject incoming ICMP ECHO request packets if their destination address is a broadcast address.

The **Syncookie Protection** field corresponds to the kernel tcp_syncookie setting. Setting Syncookie Protection to "yes" will cause the kernel to protect against SYN flood attacks.

Dummy Interface

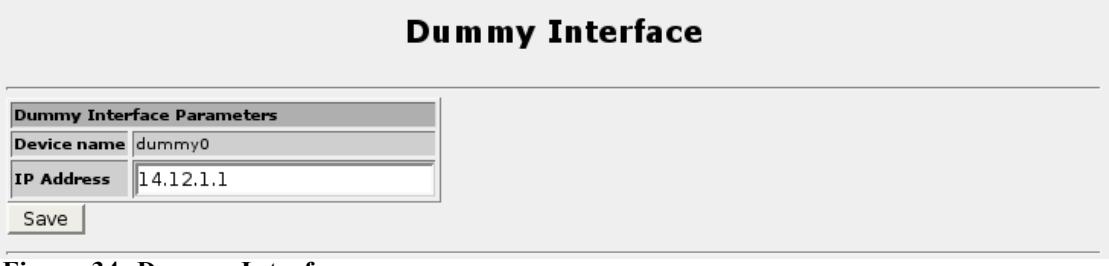

Figure 34: Dummy Interface

This menu allows you to configure a dummy interface. Normally the router is reachable on any of its interface addresses, whether the interface is active or not. When OSPF and link detection is used, inactive interfaces are not advertised to the network and thus not reachable. A dummy interface is always advertised and thus reachable.

Routing And Gateways

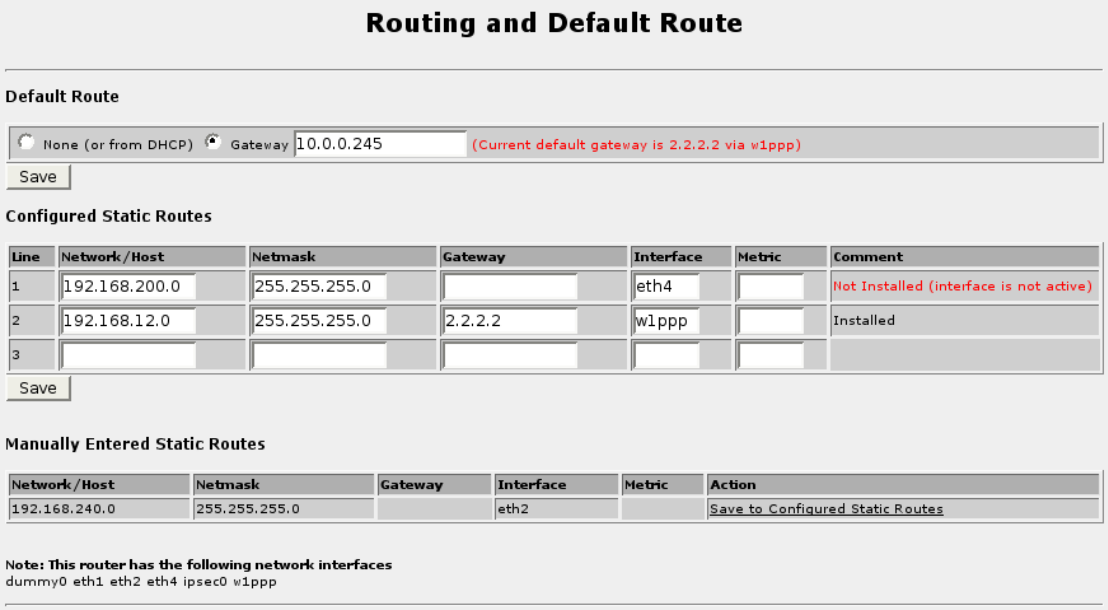

Figure 35: Routing And Gateways

This menu allows you to configure the default gateway address and static routes. Static routes specify a way to forward subnets of traffic that cannot be associated with the subnets of configured interfaces. The gateway address is the address that is used to forward traffic that can not be routed to configured interfaces or to static routes.

This menu also allows user to convert manually entered static routes to permanently configured static routes.

Default Route Table

The first table of this menu configures the default gateway address.

Note: *Don't configure a default gateway if you plan to provide one from a WAN, PPPoE or modem interface. Don't manually configure configure the default gateway in the /etc/network/interfaces file, configure the default gateway from this menu.*

If the default gateway is configured but the actual default gateway in use is different, the menu will display a warning accompanied by the actual gateway. Use the **Save** button below the table to change the default gateway setting.

Configured Static Routes

This table configures static and host routes.

The **Network/Host** and **Netmask** fields describe the remote network the static route will reach. If the netmask field is not entered (or a netmask of 255.255.255 is entered) the routing will define a host route. Any other netmask will define a network route. If the network field is cleared the route will be deleted upon the next save.

The **Gateway** field describes an address that is used as the next hop to forward traffic to. If this field is not specified than traffic is forwarded to the Interface.

The **Interface** field describes the network interface this static route will use. The interface does not need to be active or even exist, but the route will not be installed until both are true. You do not need to provide an interface, but doing so will cause the menu to warn you if the gateway is not owned by the interface. The menu provides a list of currently configured interfaces for your convenience.

The **Metric** field specifies an integer cost metric for the route, which is used when choosing among multiple routes in the routing table that most closely match the destination address of a packet being forwarded. The route with the lowest metric is chosen.

The **Comment** field shows the status of the static route, and provides a basic cause when the route is not installable.

The **Save** button below the table will save the routes and immediately install them. The following sanity checks will be made for static routes:

- The Netmask can not be 0.0.0.0.
- If the interface is active the static route will be installed, if it can not be installed, it will be treated as illegal.
- A routings Gateway address must be owned by the routings interface.

Delete routes by removing their Network/Host addresses before saving.

Manually Entered Static Routes

This table will be shown if there are active static routes which are not in the Configured Static Routes table. Following a routes "Save to Configured Static Routes" link will make the route permanent.

Note: *There are situations where manually entered routes should not be converted, e.g. routes dynamically added by IPsec and GRE tunnels. Making these routes permanent may cause the daemons that add them to fail.*

Static Multicast Routing

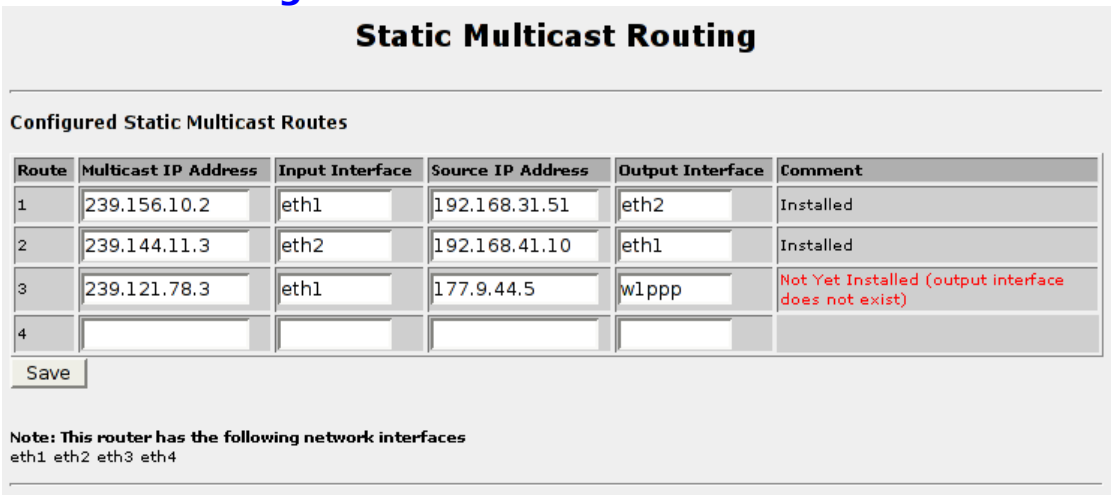

Figure 36: Static Multicast Routing

This menu allows you to configure static multicast routing.

The **Configured Static Multicast Routes** table shows configured multicast routes.

New routings may be added by completing the bottom row of the table and selecting the **Save** button. Routings may be deleted by clearing the routings **Multicast IP Address** field and selecting the **Save** button.

The **Multicast IP Address** field specifies the multicast IP address to be forwarded.

The **Input Interface** field specifies the interface upon which the multicast packet arrives.

The **Source IP Address** specifies the multicast packet's expected source IP address.

The **Output Interface** specifies the interface to which the matched multicast packet will be forwarded.

The **Comment** field shows the current status of the the routing.

The **Note** field below the table shows current active interfaces.

In order to start Multicast routing at each and every boot, you must enable it via the System folder, Bootup And Shutdown menu.

DNS Client

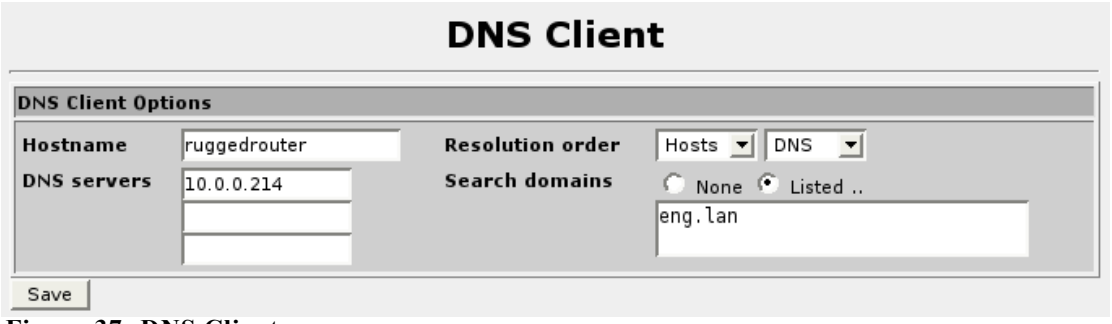

Figure 37: DNS Client

This menu allows you to display and configure various DNS client fields.

The **Resolution Order** selector determines the order of sources for resolving domain names into IP addresses. The Hosts file /etc/hosts can be populated with frequently used, but unchanging addresses. DNS refers to any configured DNS servers.

The **DNS servers** fields allow you to specify, in order, the serves to resolve from.

The **Search domains** fields allow you to specify the domain name of the network the router is located within. This allows short names relative to the local domain to be used. If you do not specify a domain name the router will try and extract this information from the host addresses.

Host Addresses

Figure 38: Host Addresses

This menu allows you to display and configure host addresses. Host addresses are useful when a non-changing IP address is often used or when DNS is not configured.

Follow the **Add a new host address** link to add an address.

End To End Backup

End To end backup is method of using two interfaces to ensure a reliable end to end connection between two routers using alternate routing, without the need to configure routing protocols.

The two interfaces are assigned as a primary:secondary backup pair. The primary interface serves as the gateway. If connectivity to the target is lost from the primary interface, traffic is migrated to the secondary interface. When connectivity is restored on the primary path, traffic will be restored to it.

Figure 39: End To End Backup Example

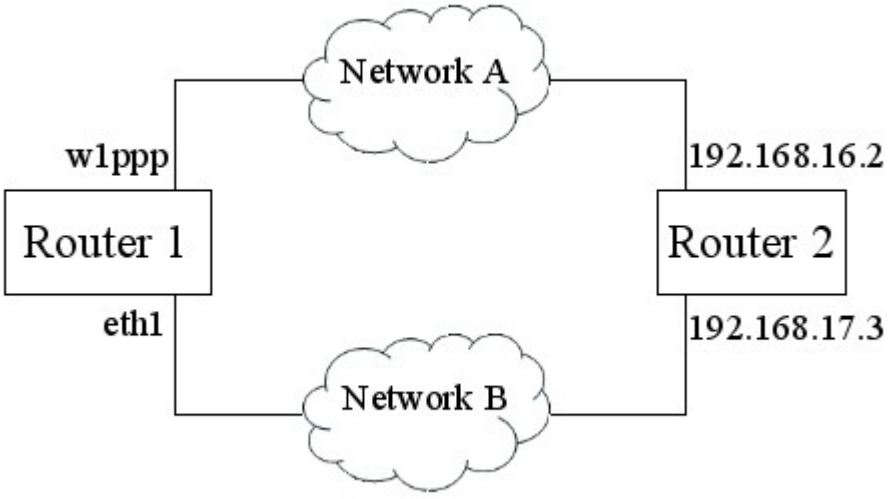

The backup is “end to end” because connectivity is determined by the availability of an interface on the target system, and not a local link. In the above figure, interface w1ppp acts as the primary interface and eth1 acts as the secondary interface. The router tests the primary path by probing 192.168.16.2 on router 2. A failure of the either w1ppp, network A or the remote link on router2 will render the primary path as “failed”.

If the primary path fails, the routing table will be modified to direct packets out the secondary (eth1 in the above figure).

Presumably, the secondary is a higher cost (and perhaps lower throughput) path. In the initial deployment of this feature, the secondary path was implemented with Ethernet-CDMA modem. The modem featured a low latency connection time (initiated by the reception of packets) but had a low bandwidth capability and high monetary cost.

Note that the feature must be implemented at both routers. If the feature is only implemented at router 1, the second router's gateway will still point towards Network A after a failure of the primary path. Packets from router 1 would reach router 2 through the secondary, but the responses would disappear in the black hole of the failed path.

Configuring End To End Backup

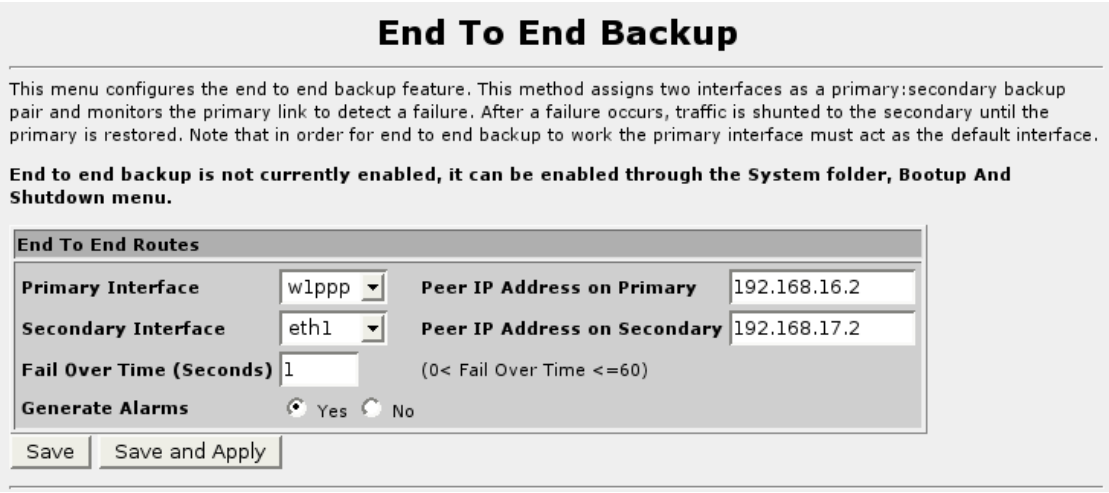

Figure 40: End To End Backup

This menu allows you to display and configure end to end backup.

In order to start end to end backup at each and every boot, you must enable it via the System folder, Bootup And Shutdown menu. The menu will remind you if the feature is not enabled.

The **Primary Interface** field determines the primary interface. The interface selected should be configured to supply the default gateway.

The **Peer IP Address on Primary** field sets the IP address to probe for connectivity on the primary interface.

The **Secondary Interface** field determines the secondary interface.

The **Peer IP Address on Secondary** field sets the IP address to probe for connectivity on the secondary interface.

The **Fail Over Timer** field determines the amount of time the primary link must be failed before directing packets down the secondary link.

The **Generate Alarms** field determines whether alarms are generated upon configuration problems and link failures.

The **Save** button will save changes to the configuration file. The **Save and Apply** button will save changes restart the end to end backup daemon.

Current Routing & Interface Table

This menu displays the current routing table and the state of the router's interfaces. Consult the **Network Utilities** chapter for details of this menu.

Chapter 5 – Configuring Ethernet Interfaces

Introduction

This chapter familiarizes the user with:

- Reading the Ethernet LEDs
- Configuring Ethernet Network Interfaces
- Configuring VLANs
- Configuring PPPoE

Ethernet Interface Fundamentals

RuggedCom manufactures dual Ethernet Interface boards in a variety of formats. Some (most notably the optical interfaces) have the same outward appearance but different order numbers. A complete set of descriptions is displayed on the console during boot and can be found after boot in the file /var/cache/ruggedrouter/inventory.

LED Designations

The RuggedRouter includes two sources of LED indicated information about Ethernet ports, the front panel LEDs and the LED Panel.

A LED is associated with each port, next to the Ethernet interface RJ45 socket. This LED is off when the link is disconnected, remains solidly on when the link is established and flashes briefly from on to off when traffic occurs.

The LED Panel also summarizes this information. LEDs 1-4 reflect traffic on Ethernet port 1-4. LEDs 5-8 reflect the link status of the same ports.

VLAN Interface Fundamentals

A virtual LAN (VLAN) is a group of devices on one or more LAN segments that communicate as if they were attached to the same physical LAN segment. VLANs are extremely flexible because they are based on logical instead of physical connections. When VLANs are introduced, all traffic in the network must belong to one or another VLAN. Traffic on one VLAN cannot pass to another, except through an intranetwork router or layer 3 switch.

The IEEE 802.1Q protocol specifies how traffic on a single physical network can be partitioned into VLANs by “tagging” each frame or packet with extra bytes to denote which virtual network the packet belongs to.

VLAN Tag

A VLAN tag is the identification information that is present in frames in order to support VLAN operation. The 4-byte VLAN tag is inserted into the Ethernet frame between the Source MAC Address field and the Length/Type field. The first 2-bytes of the VLAN tag consist of the "802.1Q Tag Type" and are always set to a value of 0x8100.

The last 2-bytes of the VLAN tag contain the following information: the first 3-bits are a User Priority Field that may be used to assign a priority level to the Ethernet frame. The next 1-bit is a Canonical Format Indicator (CFI) used in Ethernet frames to indicate the presence of a Routing Information Field (RIF). The last 12-bits are the VLAN Identifier (VID) which uniquely identifies the VLAN to which the Ethernet frame belongs.

RuggedRouter Functions Supporting VLANs

Functions	***Supported ?***	***Comments***
Static Route and Default Route	Y	
Static Multicast Routing	Y	
End To End backup	Y	
PPPoE	N	
Shorewall Firewall	Y	
IPSec	N	Netkey (policy based VPNs) supports VLAN Klips (route based VLANs) do not support VLAN
VRRP	Y	
Traffic Prioritization	Y	
Dynamic Routing		Both OSPF and RIP support VLAN
GRE Tunnel	Y	
DHCP Server	Y	

PPPoE On Native Ethernet Interfaces Fundamentals

The RuggedRouter supports PPPoE (Point-to-Point Protocol Over Ethernet) over both external modems (described here) and internal interfaces (described in the chapter “PPPOE On ADSL”). The PPPOE On ADSL chapter contains more useful information on PPPOE Authentication, Addresses, DNS Servers and MTU Issues.

Only one PPPoE interface can be created on each Ethernet Interface. Each PPPoE interface name is assigned internally. The name is “pppX”, where X is 10 plus the native Ethernet interface the PPPoE is created upon (e.g. a PPPoE on eth1 is ppp11).

Ethernet

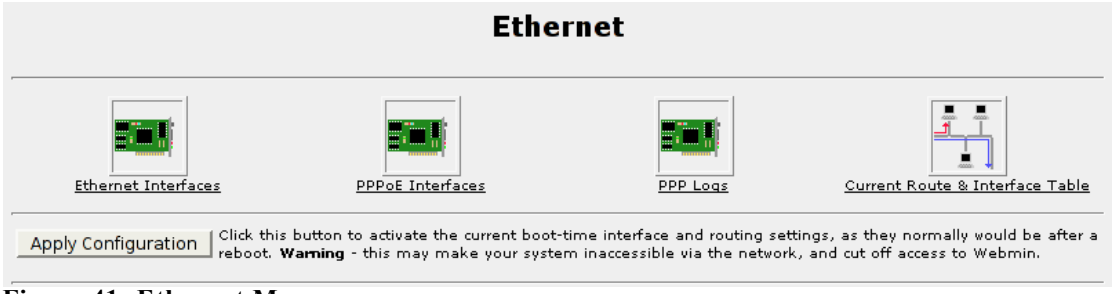

Figure 41: Ethernet Menu

This menu allows you to configure Ethernet interface parameters as well as display the routes and status of all network interfaces.

Select the **Ethernet Interfaces** icon to configure Ethernet interfaces.

The Network Interfaces menu lets you edit the permanent configuration of Ethernet interfaces, or simply try out changes. The **Apply Configuration** button serves to restore the permanently saved changes and restart Ethernet networking.

Ethernet Interfaces

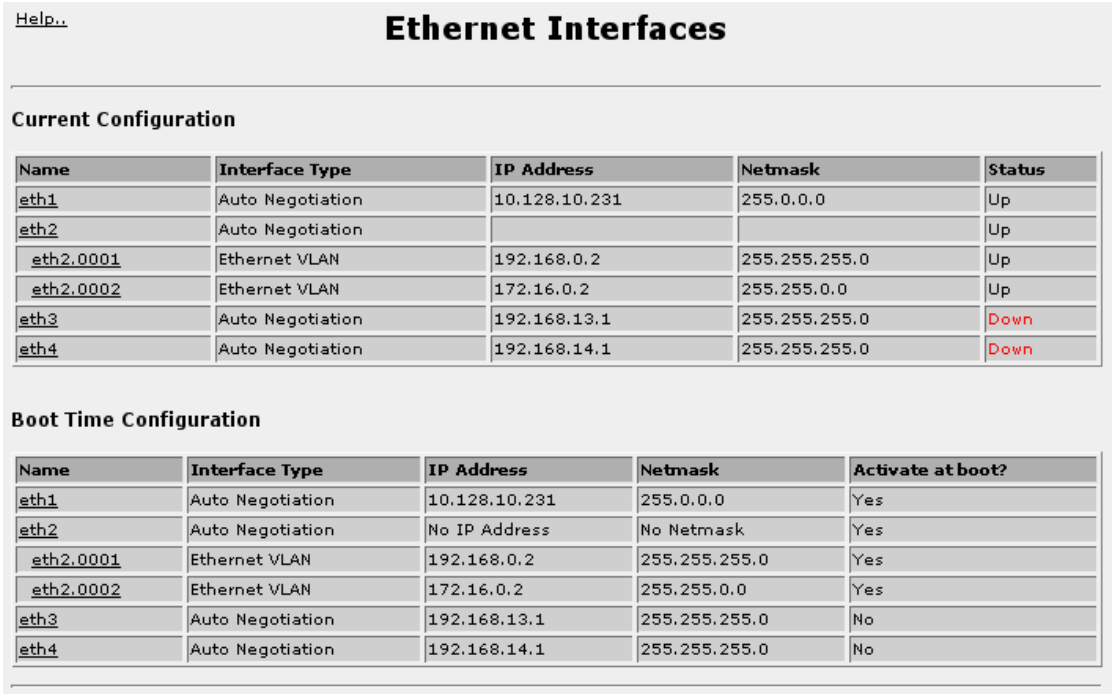

Figure 42: Current and Boot Time Ethernet Configuration

This menu allows you to display and configure the Ethernet interfaces in the router.

The **Current Configuration** table allows you to try out changes on the existing interfaces before making permanent changes. Any changes made take effect immediately, but will not be present after the next boot. The entries in this table can also be used to temporarily disable or re-enable an interface.

The **Boot Time Configuration** table router allows you make changes to the "permanent" configuration of any interface.

The Network Configuration menu **Apply Configuration** button applies permanent changes and restart Ethernet networking. If only temporary changes have been made, the permanent configuration will be re-applied.

In either table, edit the desired interface by clicking on its link under the **Name** column.

Editing Currently Active Interfaces

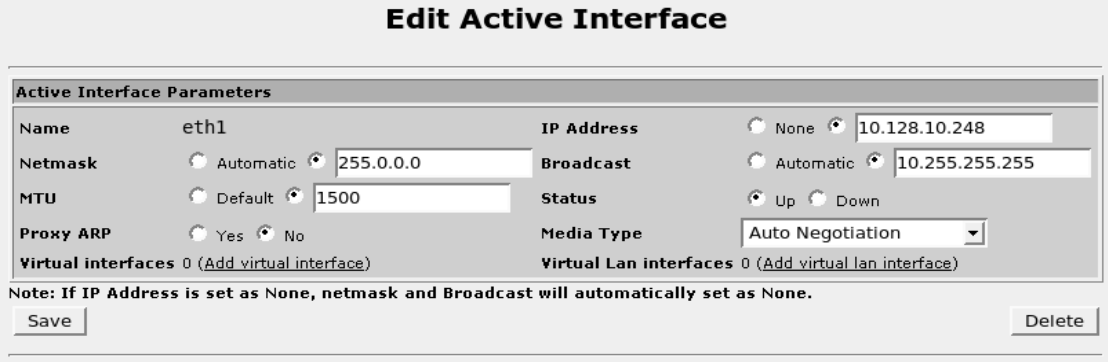

Figure 43: Editing a Network Interface

This menu allows you to make changes to the currently active interfaces. The Save button **will activate any changes, and will not affect the permanent configuration**.

The **IP Address** field sets the IP address for this interface.

The **Netmask** fields set the IP network mask for this interface. Setting this to Automatic causes the mask to be set to the usual class A, B or C network mask (as derived from the interface address. The next field can be used to specify the mask manually.

The **Broadcast** fields set the IP broadcast address for this interface. Setting this to Automatic causes the address to be set to the usual address (as derived from the interface address. The next field can be used to specify the broadcast address manually.

The **MTU** fields sets the Maximum Transfer Unit of an interface. This limits the maximum size of frames on the interface.

The **Status** field provides a way to disable the interface or bring it back into service.

The **MAC address** field displays the current Media Access Control address and allows it to be modified.

The **Proxy ARP** fields display whether the interface has proxy-arp activated.

The **Media Type** field displays the current media type. Copper interfaces may be configured to Auto-negotiable, 10 BaseT Half Duplex, 10 BaseT Full Duplex, 100 BaseT Half Duplex and 100 BaseT Full Duplex modes.

Virtual Interfaces

Use virtual interfaces when you have an Ethernet port that has multiple "real" IP addresses assigned to it, e.g. as with a port provided by an an Internet Service Provider.

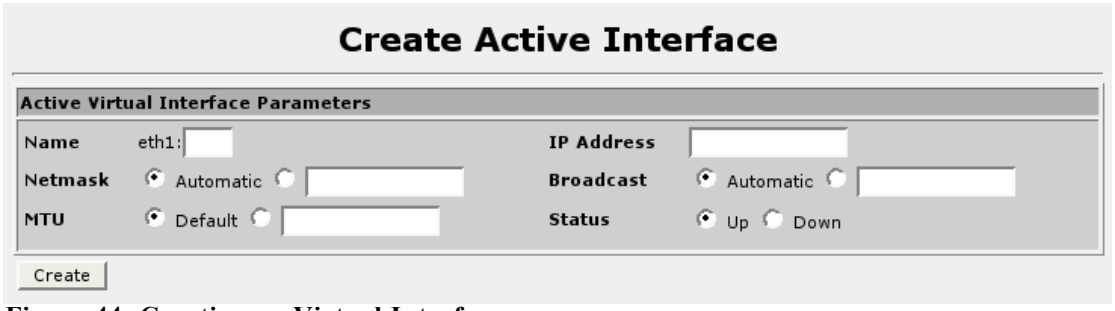

Figure 44: Creating an Virtual Interface

The only new parameter is the virtual interface descriptor, which must be a numeric value. As an example a virtual interface numbered 0 on eth1 appears as eth1:0 in interface descriptions and routing tables.

Virtual Lan Interfaces

Click the link “Add Virtual Lan Interface” when you want to create a VLAN interface.

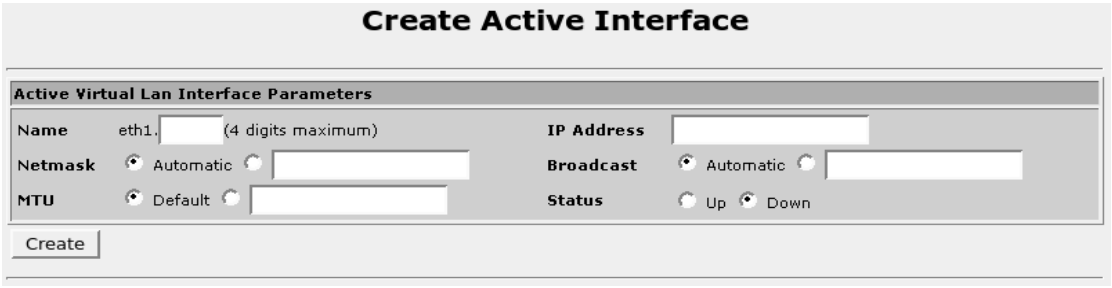

Figure 45: Creating an Virtual Lan Interface

The only new parameter is the vlan id, which must be a numeric value between 1 and 4094. The vlan id will be changed automatically as 4 digits (prefixed with 0) if the input is less than 4 digits. For example, if the input is 2, it will be automatically changed to 0002.

Edit Boot Time Interfaces

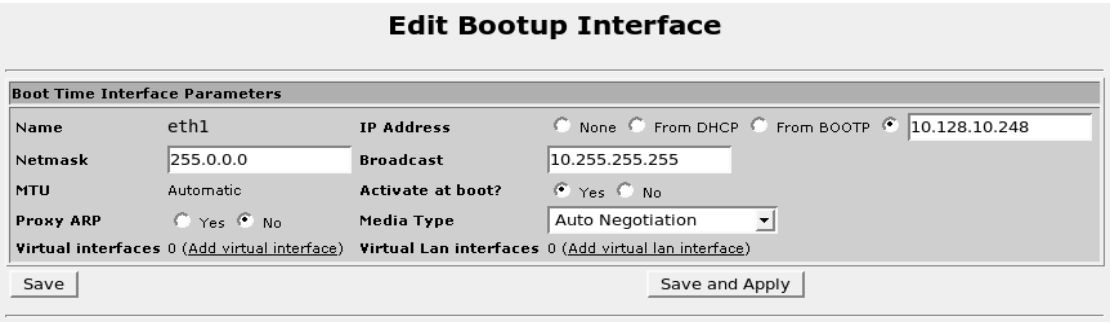

Figure 46: Editing a Boot Time Interface

This menu allows you to make permanent changes to interfaces and to immediately apply those changes if desired. The Save button will **save changes to the permanent configuration**.

The **Netmask**, **Broadcast**, **MTU, Virtual Interfaces, Proxy ARP** and **Media Type** controls are as described above.

The **IP Address** fields allow you to manually specify an IP address for this interface, or to obtain the address from DHCP or from BOOTP.

The **Activate at boot** fields allow you permanently disable the interface without actually deleting it.

The **Save and Apply** button applies any changes after they have been saved.

The **Delete and Apply** button deletes both the boot time and active interface.

The **Delete** button deletes the boot time interface but leaves the active interface in existence.

PPPoE On Native Ethernet Interfaces

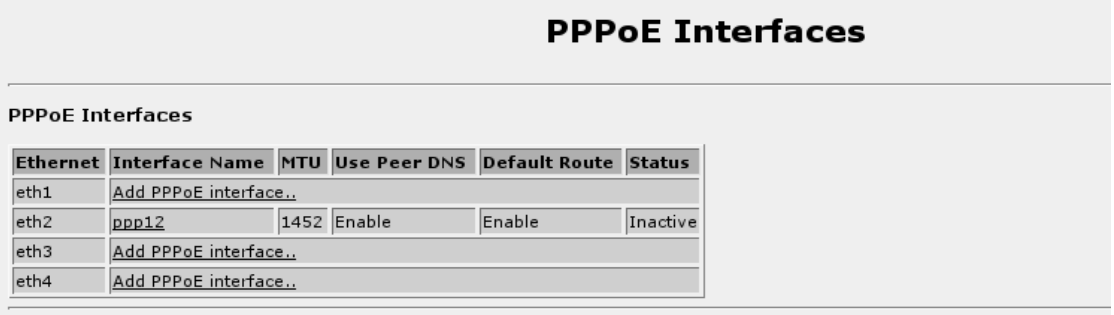

Figure 47: List PPPoE Interfaces

This menu allows you to display and configure the PPPoE interfaces on all available Ethernet ports.

The PPPoE Interfaces table allows you to add a PPPoE interface on an Ethernet ports or change PPPoE interface parameters of created interfaces. Only one PPPoE interface can be created on each Ethernet port.

The **Ethernet** field shows all available Ethernet ports.

The **Interface Name** field shows created PPPoE interfaces and provides a link to edit the existing configuration or create a new one.

The **MTU**, **Use Peer DNS** and **Default Route** fields are the configured information for PPPoE interfaces.

The **Status** field shows the current PPPoE link status.

Edit PPPoE Interface

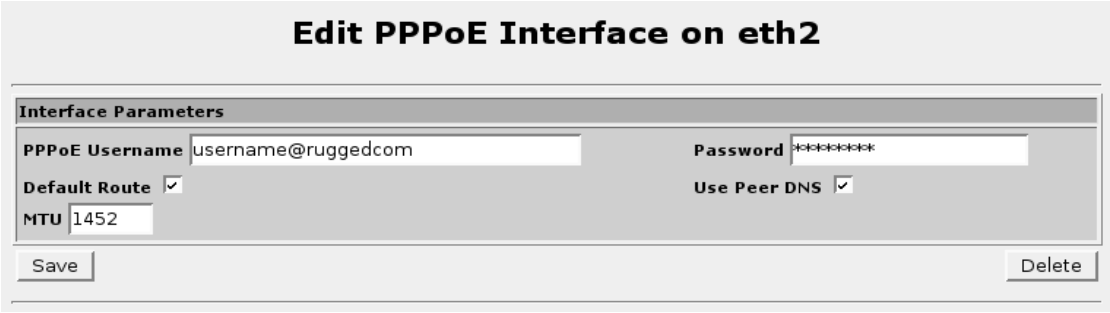

Figure 48: Editing a PPPoE Interface

This menus allows you to edit a PPPoE interface.

The **PPPoE Username** field determines the username to use when connecting to the PPPoE server as specified by your provider.

The **Password** field determines the password provided to the PPPoE server.

The **Default Route** checkbox enables automatically setting a default route using this interface whenever it connects. If this is your primary connection you probably want this option enabled.

The **Use peer DNS** checkbox enables automatically setting the DNS server entries that the PPPoE server recommends. Enable this option unless you provide your own name servers.

The **MTU** field defines the MTU size to request when connecting to the PPPoE server. In some cases the PPPoE provider may provide a smaller MTU in which case the smaller setting will be used, or it may refuse to alter the MTU and use whatever it considers to be the default.

The **Save** button will update all of the changes. The current PPPoE link will be connected.

The **Delete** button will delete the PPPoE interface, closing the current PPPoE link.

PPP Logs

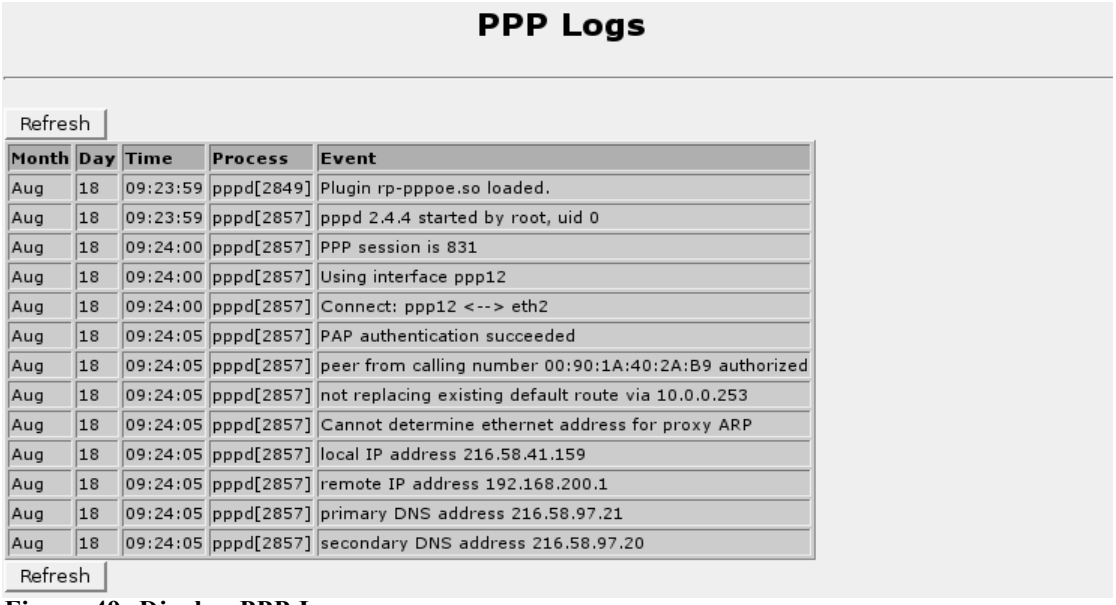

Figure 49: Display PPP Logs

This menu displays the native Ethernet and internal ADSL interface PPPoE connection messages. This is mainly useful when trying to debug a PPP connection problem.

Current Routes & Interface Table

The table provided by this command is as described in the **Networking** menu, **Network Utilities** sub-menu. It is also provided here as a convenience.

Chapter 6 – Configuring Frame Relay/PPP And T1/E1

Introduction

This chapter familiarizes the user with:

- Frame Relay and PPP Terminology and Issues
- Configuring Frame Relay and PPP Links
- Viewing status and statistics
- Upgrading Firmware

T1/E1 Fundamentals

A T1 is a communications circuit upon which has been imposed a digital signal 1 (DS1) signaling scheme. The scheme allows 24 "timeslots" of 64 Kbps DS0 information (as well as 8 Kbps of signaling information) to be multiplexed to a 1544 Kbps circuit.

The 24 DS0s can be used individually as standalone channels, bonded into groups of channels or can be bonded to form a single 1536 Kbps channel, referred to as a clear channel. Not all channels need be used. It is quite common to purchase N channels of 64Kbps bandwidth and leave the remainder unused, this is known as fractional T1.

The telephone network terminates the T1 line and maps each of the channels through the T1 network to a chosen T1 line. Individual and bonded DS0s from more than one remote T1 can be aggregated into a full T1 line (often referred to as central site concentration).

Whereas the T1 line itself is referred to as the physical interface, groups of DS0s form channels and the protocols that run on the channels are known as a logical interfaces. The RuggedRouter provides you the ability to operate Frame Relay or PPP over your logical interfaces.

An E1 is is a communications circuit conforming to European standards, possessing 32 64 Kbps channels, of which one is usually reserved for signaling information.

Frame Relay

Frame Relay is a packet switching protocol for use over the WAN. The RuggedRouter provides the ability to construct point-to-point IP network connections over Frame Relay.

Each Frame Relay interface provides a "link" between a local and peer station. One of the stations must be configured as a Data Communications Equipment (DCE) device (often known as the "Switch") while the peer station must be configured as a Data Terminal Equipment (DTE) device (often known as Customer Premises Equipment (CPE)). The DCE is responsible for managing the link, advertising connections to the DTE and switching packets between connections. The DTE raises individual connections and sends data on them.

When using a T1/E1 line to access a public Frame Relay provider, configure the Router as a DTE.

Unlike PPP, a Frame Relay link can provide multiple (up to 990) connections. Each connection is identified by a Data Link Connection Identifier (DLCI) and must match at the DCE and DTE. The use of multiple connections can support meshed network interconnections and disaster recovery.

Location Of Interfaces And Labeling

Unlike the Ethernet ports (which are statically located), the location of T1/E1, DDS and ADSL ports in your router depends upon the number of ports and how they were ordered. Refer to the labeled hardware image as presented in the Webmin home page.

To make labeling easy to understand, all T1E1, T3, DDS and ADSL ports are assigned a unique port number that relates to the LEDs on the status panel.

LED Designations

The RuggedRouter includes two sources of LED indicated information about T1E1 lines, the T1E1 card itself and the LED Panel.

One LED is associated with each line, next to the interface jack. This LED is red when the link is disconnected, flashes green when the link is connecting and remains solid green when the link is established.

The RuggedRouter also indicates information about T1E1 ports on the LED Panel. A pair of LEDs will indicate traffic and link status of the port. Consult the section "Using The LED Status Panel" to determine which LEDs correspond to the port.

Included With T1E1

T1E1 includes wanpipemon, a utility that can capture traces from the T1E1 line.

T1/E1

Figure 50: T1/E1 Trunks And Interfaces

This menu allows you to display and configure T1 or E1 Trunks as well as display the routes and status of the network interfaces.

T1/E1 Network Interfaces

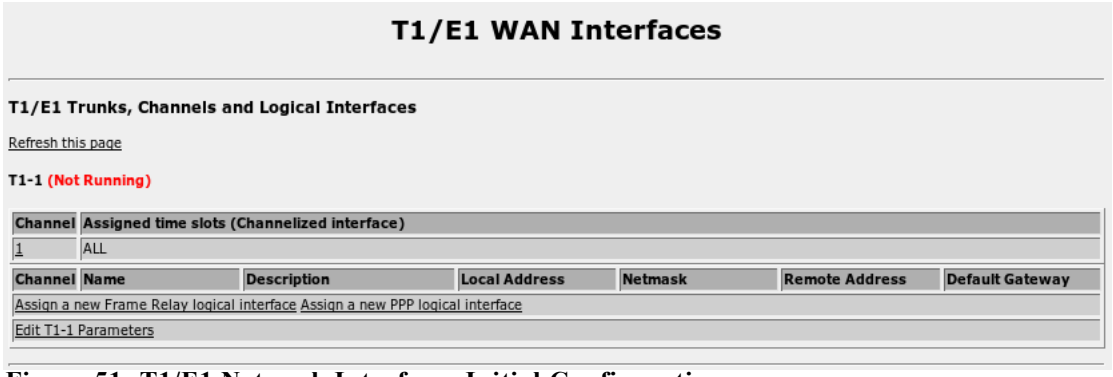

Figure 51: T1/E1 Network Interfaces Initial Configuration

This menu allows you to display and configure T1/E1 Trunk parameters, Channels and the logical interfaces that run on them. A table is presented for each interface.

Note that the interface number is the same regardless of whether it is a T1 or E1 interface. Interface numbers are as described by the "WAN" labels as shown in the home page chassis diagram.

The status of the trunks physical and logical interfaces are shown This menu presents connection statuses but does not update them in real time. Click on the **Refresh this page** link to update to the current status.

Strategy For Creating Interfaces

Initially, each interface will be configured as T1 and will have a single channel that includes all timeslots (1-24). Channelized cards can have their timeslots reassigned to make additional channels. Unchannelized cards may have timeslots removed from their single timeslot.

If the interface is to be an E1, convert it using the "Edit T1-1 Parameters" link.

If the interface is channelized and you need to have more than one channel, construct the channel groups with the desired bandwidths. This can be done by editing the single initially configured channel and removing timeslots. The unassigned timeslots will be displayed on the main menu in a link that creates channels, as shown below.

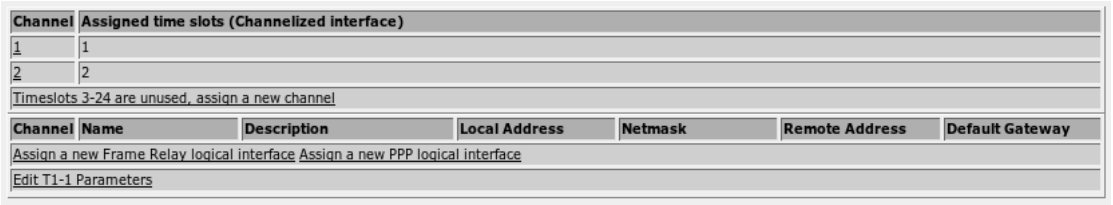

Figure 52: T1/E1 Network Interfaces After Channel Creation

Once all timeslots have been assigned to channels, the "Timeslots.." link will no longer appear. Note that you do not have to assign all timeslots.

Assign Frame Relay or PPP to the channels by following the "Assign .. Protocol" links. The resultant menus will allow you select the desired channel.

If you are assigning multiple DLCIs, assign the first DLCI used by that interface and configure the Frame Relay Link Parameters and that DLCIs network parameters.

After assigning the first DLCI, you may revisit the interface through the link under the **Name** field and add additional DLCIs.

Once all channels have been assigned, the "Assign" links will no longer appear, as shown below. Note that any of the Frame Relay interfaces on a channel (in this case w1c4fr16 and w1c4fr17) may be used to edit the Frame Relay Link Parameters.

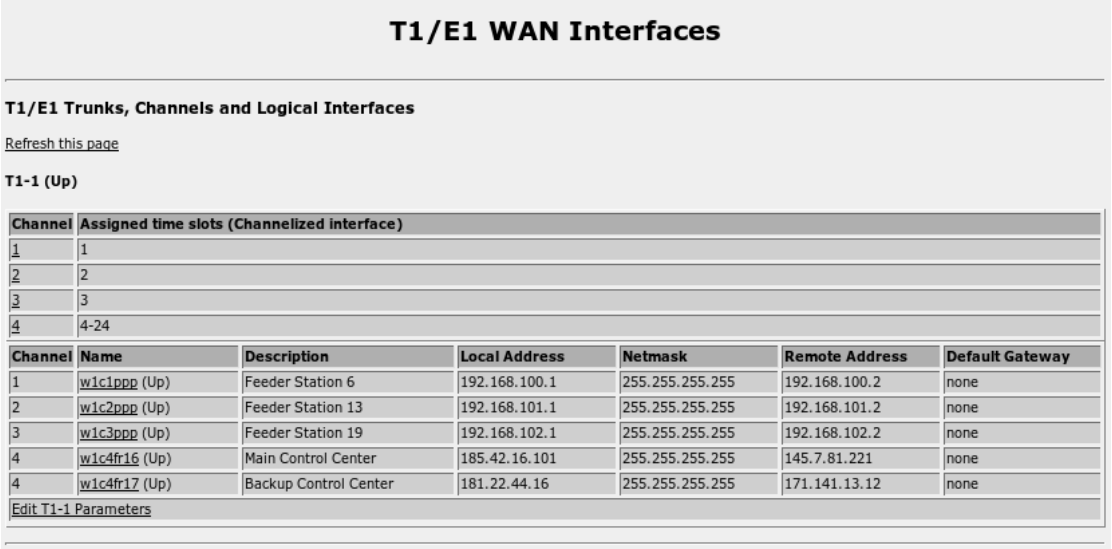

T1/E1 WAN Interfaces

T1/E1 Trunks, Channels and Logical Interfaces

Refresh this page

T1-1 (Up)

Channel	Assigned time slots (Channelized interface)
1	1
2	2
3	3
4	4-24

Channel	Name	Description	Local Address	Netmask	Remote Address	Default Gateway
1	w1c1ppp (Up)	Feeder Station 6	192.168.100.1	255.255.255.255	192.168.100.2	none
2	w1c2ppp (Up)	Feeder Station 13	192.168.101.1	255.255.255.255	192.168.101.2	none
3	w1c3ppp (Up)	Feeder Station 19	192.168.102.1	255.255.255.255	192.168.102.2	none
4	w1c4fr16 (Up)	Main Control Center	185.42.16.101	255.255.255.255	145.7.81.221	none
4	w1c4fr17 (Up)	Backup Control Center	181.22.44.16	255.255.255.255	171.141.13.12	none

Edit T1-1 Parameters

Figure 53: T1/E1 Network Interfaces After Interface Creation

Naming Of Logical Interfaces

Webmin names the logical interfaces for you (but allows you to provide a description). All interfaces start with a "w" to identify them as wan interfaces, followed by the physical interface number.

Unchannelized hardware interfaces supply only one channel (that can be composed of a varying number of timeslots) logical interface. You may configure one PPP interface or up to 990 Frame Relay DLCI interfaces. The next part of the identifier is either "ppp" or "frX" where X the frame relay channel number.

Channelized hardware allows more than one logical interface. The next part of the identifier indicates the channel the interface uses with a "c" followed by the lowest channel used. The final part of the identifier is either "ppp" or "fr" and the frame relay channel number.

Note: *Once a channel is created, and an interface is constructed on it, the name of the interface will never change. This will remain true even if the number of timeslots on the channel is changed. This property is desirable since interface names used by features such as OSPF, RIP and the firewall can rely on the interface name. Channel re-assignments can, however, lead to a non-intuitive relationship between channels and timeslots.*

Editing A T1/E1 Interface

Figure 54: Edit T1 Interface

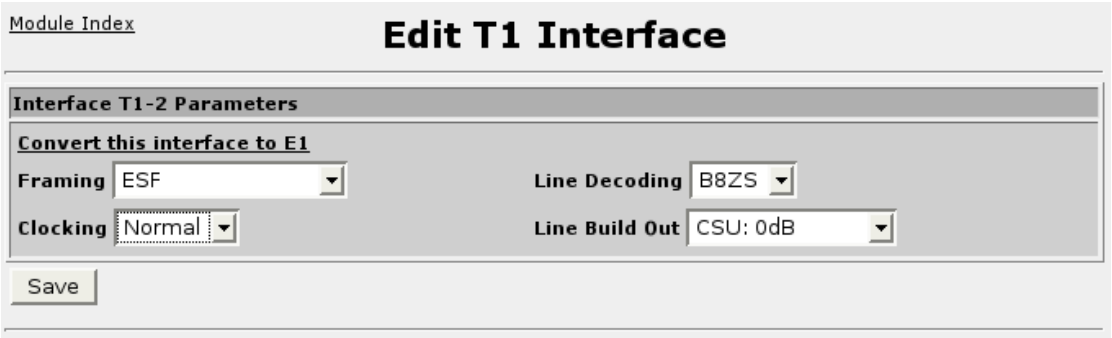

This menu allows you to display and configure T1 or E1 Trunk parameters. By default the interface is set for T1 operation. The **Convert this interface to E1** link will set the interface for E1 operation and allow you to configure its settings.

If logical interfaces use a channel above 24 and an attempt to convert from E1 to T1 will prompt to delete the logical interface first.

T1 Settings

The **Framing** field determines the framing format used. Your line provider will indicate the correct format. Modern facilities usually employ Extended Super Frame (ESF), an enhanced T1 format that allows a line to be monitored during normal operation.

The **Line Decoding** field reflects the line encoding/decoding scheme. Almost all T1s now use B8ZS.

The **Clocking** field selects whether to accept or provide clocks. In normal use the central office provides clocks and your setting should be "Normal". You may also connect to another router by using a cross-over cable and selecting a "Master" clocking option on one of the two routers.

The **Line Build Out** field "tunes" the shape of the T1 pulses and adjusts their amplitude depending upon distances and the desired attenuation.

E1 Settings

The **Framing** and **Line Decoding** fields for E1 reflect the European variants.

The **Clocking** field performs the same function as that described for T1.

Editing A Logical Interface (Frame Relay)

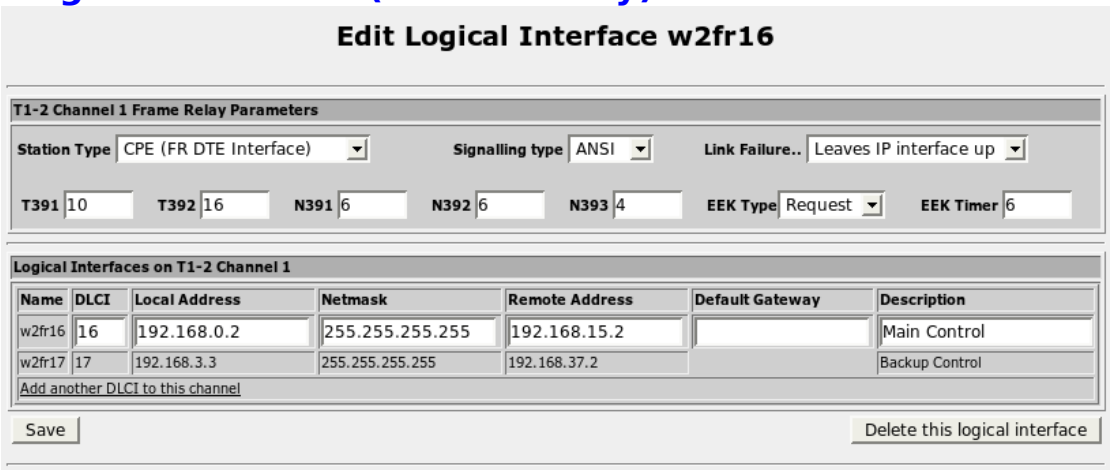

Figure 55: Edit Logical Interface (Frame Relay)

This menu allows you to configure Frame Relay link and logical interface fields.

Frame Relay Link Parameters

The first table presents the link parameters and applies to all logical interfaces.

The **Station Type** field determines whether the router acts as a customer premises equipment or as a frame relay switch. When a Frame Relay network provider is used, the CPE interface should be chosen. When the connection is end to end, it is typical to set the central site end to switch and the remote end to be CPE.

The **Signaling type** field reflects the Frame Relay link management protocol used, which include ANSI T1.617 Annex D, LMI and Q.933 signaling.

The **Link Failure** field determines whether the IP interface should reflect the state of the T1 (connected/disconnected). If you are using SNMP, enable this option as SNMP uses the state of the interface to determine the state of the connection.

The **T391** (Link Integrity Verification polling) timer is valid at the CPE and indicates the number of seconds between the transmission of In-channel Signaling messages.

The **T392** (verification of polling cycle) timer is valid at the Switch and indicates the expected number of seconds between the reception of In-channel Signaling messages transmitted by the CPE.

The **N391** counter is valid at the CPE and defines the frequency of transmission of Full Status enquiry messages.

The **N392** counter is valid at both the CPE and the Switch and defines the number of errors during N393 events which cause the channel to be inactive.

The **N393** counter is valid at both the CPE and the Switch and is an event counter for measuring N392.

The **EEK Type** field controls whether End to End Keepalive messages are sent while operating as a CPE device. If this option is set to “Off”, EEK is disabled. If this option is set to “Request”, EEK messages are sent every **EEK Timer** x T391 seconds. This timer may be configured from 1 to 100 periods in duration.

Your network provider will inform you of what is proper for these parameters.

Frame Relay DLCIs

The second table provides a listing of all DLCIs available on the channel. Only the DLCI selected from the main menu can be edited, although another DLCI can be added by following the **Add another DLCI to this channel** link.

The **DLCI Number** refers to the Data Link Connection Identifier. This number should be provided to you by your provider.

The **Local IP Address** field defines the IP address for this interface.

The **Netmask** field defines the network address mask The value 255.255.255.255 specifies a point-to-point connection which is almost always correct.

The **Remote IP Address** field defines the IP address for other side of this interface. As most WAN links are of point-to-point type, there is only one host connected to the other end of the link and its address is known in advance. This option is the address of the 'other end' of the link and is usually assigned by the network administrator or Internet service provider.

The **Use as Default Route** fields allow you to install a default route to be used while the interface is active. If specified, the gateway address should reside within the host portion of the subnetted remote IP address.

The **Description** field attaches a description to the logical interface viewable from the network interfaces menu.

The **Delete this logical interface** button removes the currently selected interface. Repetitive use of this button on other DLCIs assigned to the channel will free the channel up.

Editing A Logical Interface (PPP)

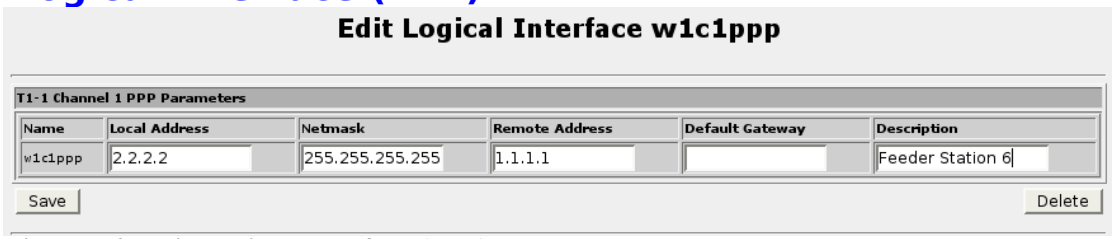

Figure 56: Edit Logical Interface (PPP)

The **Local Address**, **Netmask**, **Remote Address**, **Default Gateway and Description** fields are as described in the previous section.

T1/E1 Statistics

When at least one logical interface is configured, T1/E1 Link and logical interface statistics will be available. These statistics are available from links on the T1/E1 WAN Interfaces menu.

Link Statistics are provided through the "View Link Statistics" link at the bottom of each interface table. Frame Relay and PPP statistics are available through "(Statistics)" links under the interface name column of each interface table.

Link Statistics

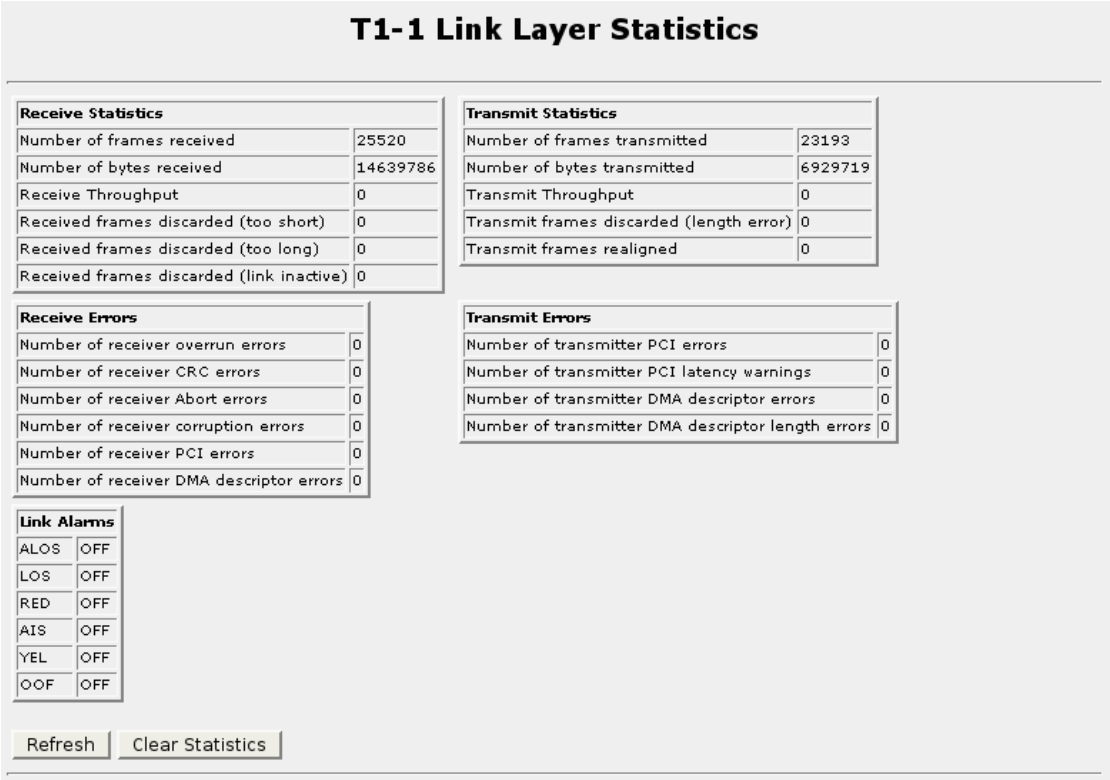

Figure 57: T1/E1 Link Statistics

The **Link Alarms** indicate ongoing problems.

ALOS/LOS (Loss of Signal) – This alarm indicates a complete absence of synchronization pulses on the line.

RED (Red Alarm) - This is a local equipment alarm. It indicates that the incoming signal has been corrupted for a number of seconds. This equipment will then begin sending a yellow alarm as its outbound signal.

AIS (Alarm Indication Signal, or BLUE alarm) - This alarm indicates the total absence of incoming signal as a series of continuous transitions (an all 1's pattern) is received.

YEL (Yellow Alarm) – This alarm is transmitted to the network and alerts it that a failure has been detected.

OOF (Out of Frame) – This alarm signifies the occurrence of a particular density of framing error events. This alarm could signify that the wrong framing mode is configured.

Frame Relay Interface Statistics

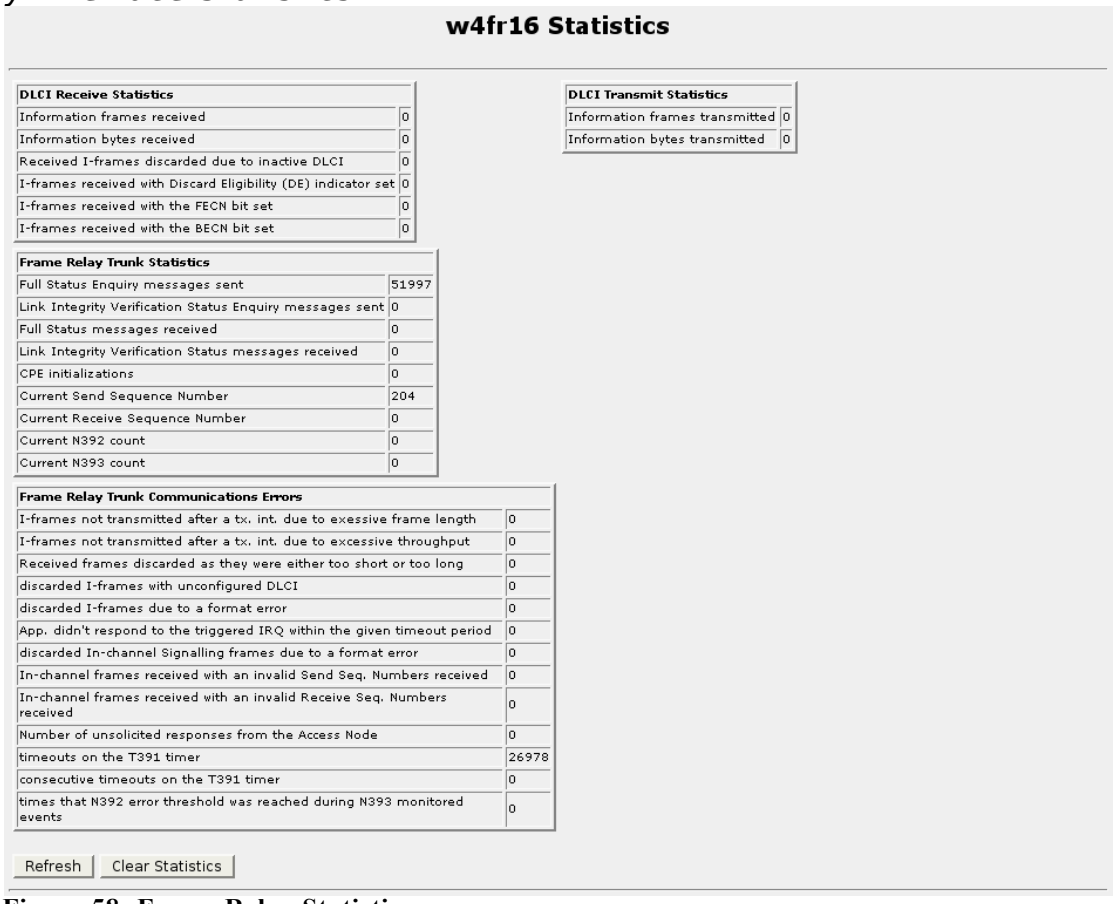

w4fr16 Statistics

DLCI Receive Statistics	
Information frames received	0
Information bytes received	0
Received I-frames discarded due to inactive DLCI	0
I-frames received with Discard Eligibility (DE) indicator set	0
I-frames received with the FECN bit set	0
I-frames received with the BECN bit set	0

DLCI Transmit Statistics	
Information frames transmitted	0
Information bytes transmitted	0

Frame Relay Trunk Statistics	
Full Status Enquiry messages sent	51997
Link Integrity Verification Status Enquiry messages sent	0
Full Status messages received	0
Link Integrity Verification Status messages received	0
CPE initializations	0
Current Send Sequence Number	204
Current Receive Sequence Number	0
Current N392 count	0
Current N393 count	0

Frame Relay Trunk Communications Errors	
I-frames not transmitted after a tx. int. due to exessive frame length	0
I-frames not transmitted after a tx. int. due to excessive throughput	0
Received frames discarded as they were either too short or too long	0
discarded I-frames with unconfigured DLCI	0
discarded I-frames due to a format error	0
App. didn't respond to the triggered IRQ within the given timeout period	0
discarded In-channel Signalling frames due to a format error	0
In-channel frames received with an invalid Send Seq. Numbers received	0
In-channel frames received with an invalid Receive Seq. Numbers received	0
Number of unsolicited responses from the Access Node	0
timeouts on the T391 timer	26978
consecutive timeouts on the T391 timer	0
times that N392 error threshold was reached during N393 monitored events	0

Refresh | Clear Statistics

Figure 58: Frame Relay Statistics

Note that the **Frame Relay Trunk Statistics** and **Frame Relay Trunk Communications Errors** tables are common to all Frame Relay DLCIs on the trunk.

PPP Interface Statistics

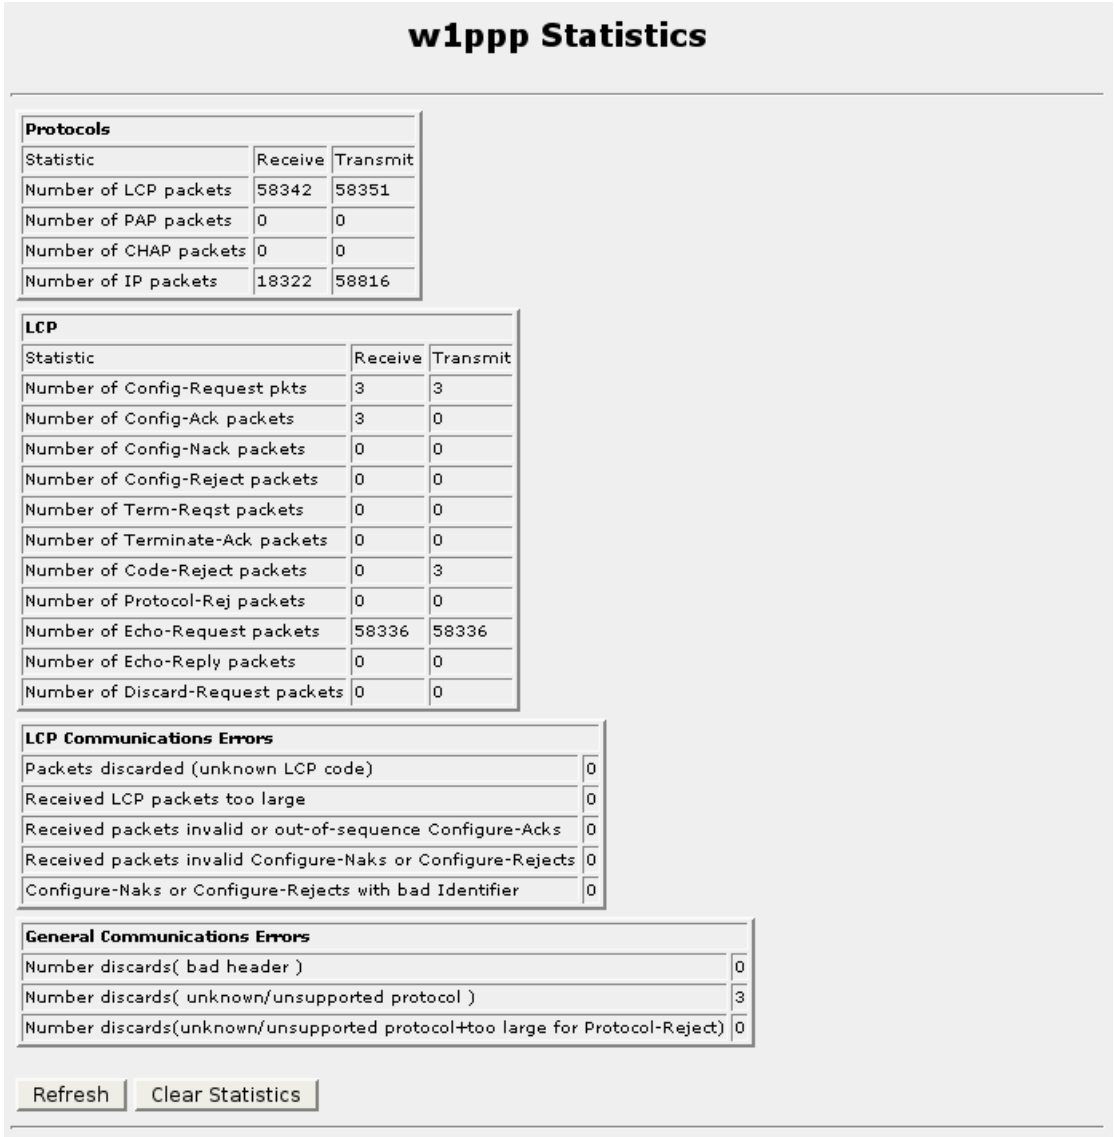

w1ppp Statistics

Protocols		
Statistic	Receive	Transmit
Number of LCP packets	58342	58351
Number of PAP packets	0	0
Number of CHAP packets	0	0
Number of IP packets	18322	58816

LCP		
Statistic	Receive	Transmit
Number of Config-Request pkts	3	3
Number of Config-Ack packets	3	0
Number of Config-Nack packets	0	0
Number of Config-Reject packets	0	0
Number of Term-Reqst packets	0	0
Number of Terminate-Ack packets	0	0
Number of Code-Reject packets	0	3
Number of Protocol-Rej packets	0	0
Number of Echo-Request packets	58336	58336
Number of Echo-Reply packets	0	0
Number of Discard-Request packets	0	0

LCP Communications Errors	
Packets discarded (unknown LCP code)	0
Received LCP packets too large	0
Received packets invalid or out-of-sequence Configure-Acks	0
Received packets invalid Configure-Naks or Configure-Rejects	0
Configure-Naks or Configure-Rejects with bad Identifier	0

General Communications Errors	
Number discards(bad header)	0
Number discards(unknown/unsupported protocol)	3
Number discards(unknown/unsupported protocol+too large for Protocol-Reject)	0

Refresh | Clear Statistics

Figure 59: PPP Link Statistics

T1/E1 Loopback

When at least one logical interface is configured, a T1/E1 Loopback tests can be performed. This menu can be reached from a link on the T1/E1 WAN Interfaces menu.

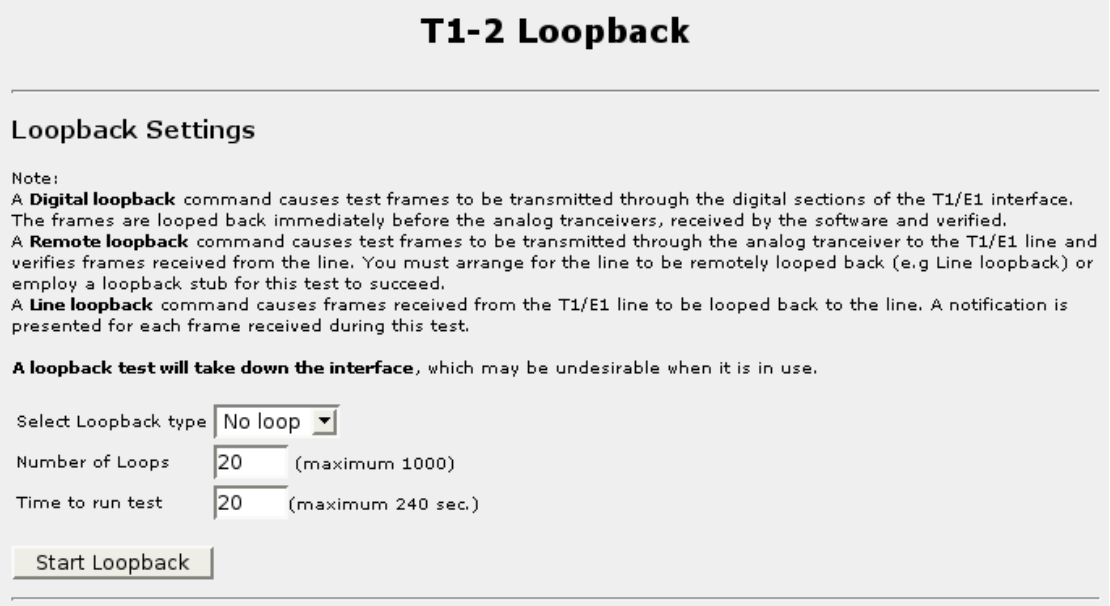

Figure 60: T1/E1 Loopback Menu

The loopback test provides a means to test the digital and analog hardware of your T1/E1 hardware and the T1/E1 line. The sender transmits a number of frames which are looped back to it. The returning frames are verified for correctness.

A digital loopback is started first, verifying the digital section of the interface. If a loopback stub is inserted in the interface jack, a remote loopback will verify the interfaces digital and analog sections. If the remote equipment is able to loop, the entire T1/E1 line can be verified. If the remote router is another RuggedCom router, a starting a line loopback will verify both cards and the line. This router will display the count of loopback frames as they arrive.

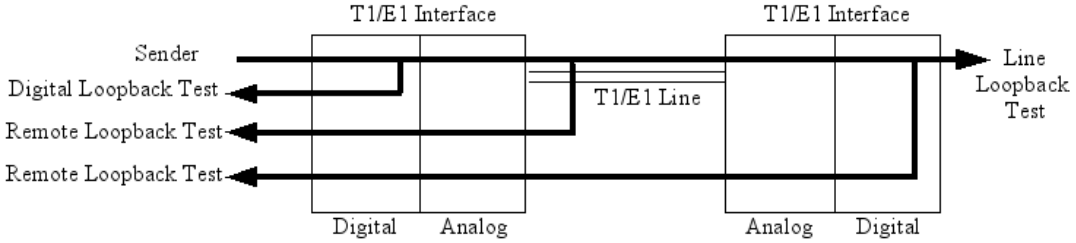

Figure 61: T1/E1 Loopback

The **Select Loopback Type** field selects the loopback.

The **Number of Loops** field controls the frames sent during digital and remote loopback. This parameter is not used during line loopback.

The **Time to run test** field limits the time the sender will transmit and the router running line loopback will wait.

Running a loop test on an active interface will immediately cause it to go down. The loop test automatically initializes the trunk after completing the test.

Current Routes & Interface Table

The table provided by this command is as described in the **Networking** menu, **Network Utilities** sub-menu. It is also provided here as a convenience.

Upgrading Software

For some customers, access to remote sites in accomplished solely by a T1 or E1 connection. Usually a software upgrade will stop the system being upgraded, perform the upgrade and then restart it. If T1E1 was upgraded in this way, the upgrade would fail as the T1E1 link was taken down. Instead, T1E1 software upgrades modify only the software on the disk. You must schedule a reboot in order to run the new version of T1E1 software.

Upgrading Firmware

RuggedCom T1E1 interfaces reside upon PCI interface cards. These cards contain FLASH memory which (from time to time) will be required to be upgraded. The upgrade process will take down the T1E1 links, upgrade the firmware and then restart the interfaces.

Note: *The upgrade process requires upwards of 15 minutes for each PCI interface card. Because of the lengthy duration required to upgrade the interfaces, RuggedCom does not automatically perform the firmware upgrade. Instead, the scheduling of the upgrade is left to the user.*

The upgrade can be performed by signing on to the platform via the console or ssh and running the command "/usr/sbin/update-wanfirmware". If the ssh connection has been made over an active T1E1 interface, the connection will fail but the upgrade will continue.

The upgrade can also be scheduled for a specific time by using the **System** menu, **Scheduled Commands** sub-menu. Set the **Commands to execute** field to "/usr/sbin/update-wanfirmware proceed", set the **Run in directory** field to "/root" and set the **Run at time** field to the desired upgrade time.

After the upgrade completes, alarms recommending an upgrade will be cleared.

Chapter 7 – Configuring Frame Relay/PPP And T3

Introduction

This chapter familiarizes the user with:

- Configuring Frame Relay and PPP Links
- Viewing status and statistics
- Upgrading Firmware

T3 Fundamentals

A T3 is a communications circuit upon which has been imposed a digital signal 3 (DS3) signaling scheme. The scheme allows 672 "timeslots" of 64 Kbps DS0 information to be multiplexed to a 44.736 Mbps circuit.

Channel groups and fractional lines are not supported.

The RuggedRouter provides you the ability to operate Frame Relay or PPP over your physical interfaces.

Location Of Interfaces And Labeling

Unlike the Ethernet ports (which are statically located), the location of T1/E1, T3, DDS and ADSL ports in your router depends upon the number of ports and how they were ordered. Refer to the labeled hardware image as presented in the Webmin home page.

To make labeling easy to understand, all T1E1, T3, DDS and ADSL ports are assigned a unique port number that relates to the LEDs on the status panel.

LED Designations

The RuggedRouter includes two sources of LED indicated information about T3 lines, the T3 card itself and the LED Panel.

One LED is associated with each line, next to the interface jack. This LED is red when the link is disconnected, flashes green when the link is connecting and remains solid green when the link is established.

The RuggedRouter also indicates information about T3 ports on the LED Panel. A pair of LEDs will indicate traffic and link status of the port. Consult the section "Using The LED Status Panel" to determine which LEDs correspond to the port.

T3 Configuration

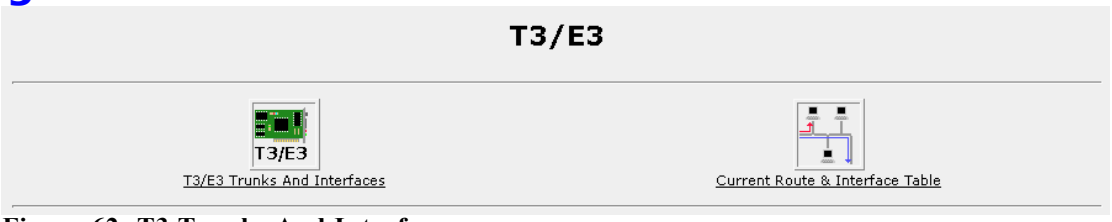

Figure 62: T3 Trunks And Interfaces

This menu allows you to display and configure T3 Trunks as well as display the routes and status of the network interfaces.

T3 Network Interfaces

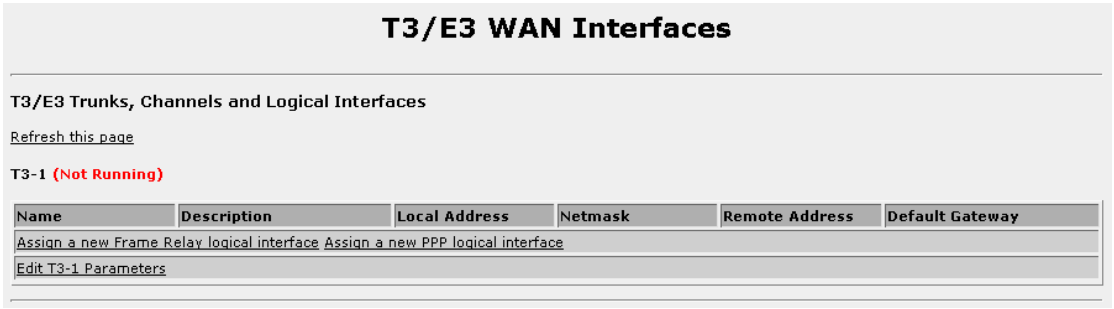

Figure 63: T3 Network Interfaces Initial Configuration

This menu allows you to display and configure T3 Trunk parameters. A table is presented for each interface.

Interface numbers are as described by the "WAN" labels as shown in the home page chassis diagram.

The status of the trunks physical and logical interfaces are shown This menu presents connection statuses but does not update them in real time. Click on the **Refresh this page** link to update to the current status.

The menu will change after assignment of a logical interface, providing links to logical interface and link statistics.

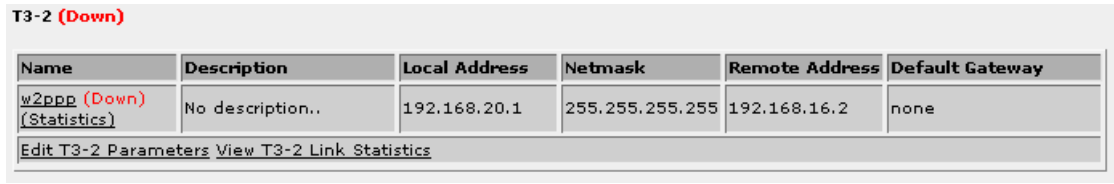

Figure 64: T3 Network Interfaces Initial Configuration

Naming Of Logical Interfaces

Webmin names the logical interfaces for you (but allows you to provide a description). All interfaces start with a "w" to identify them as wan interfaces, followed by the interface number. The next part of the identifier is either "ppp" or "fr" and the frame relay DLCI number.

Editing A T3 Interface

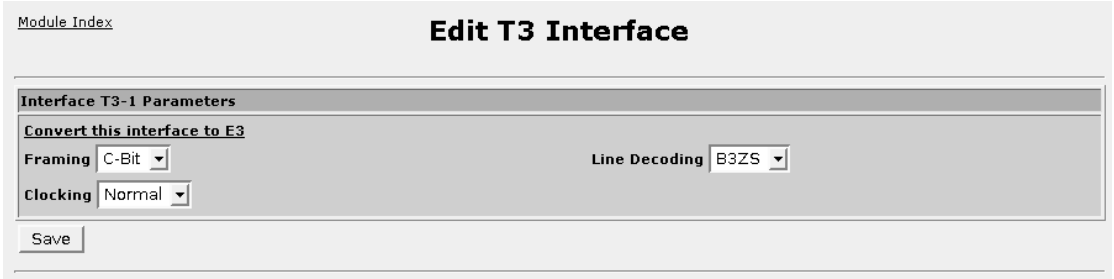

Figure 65: Edit T3 Interface

This menu allows you to display and configure T3 Trunk parameters.

The **Framing** field determines the framing format used. Your line provider will indicate the correct format.

The **Line Decoding** field reflects the line encoding/decoding scheme. Almost all T3s now use B3ZS.

The **Clocking** field selects whether to accept or provide clocks. In normal use the central office provides clocks and your setting should be "Normal". You may also connect to another router by using a cross-over cable and selecting a "Master" clocking option on one of the two routers.

Editing A Logical Interface (Frame Relay)

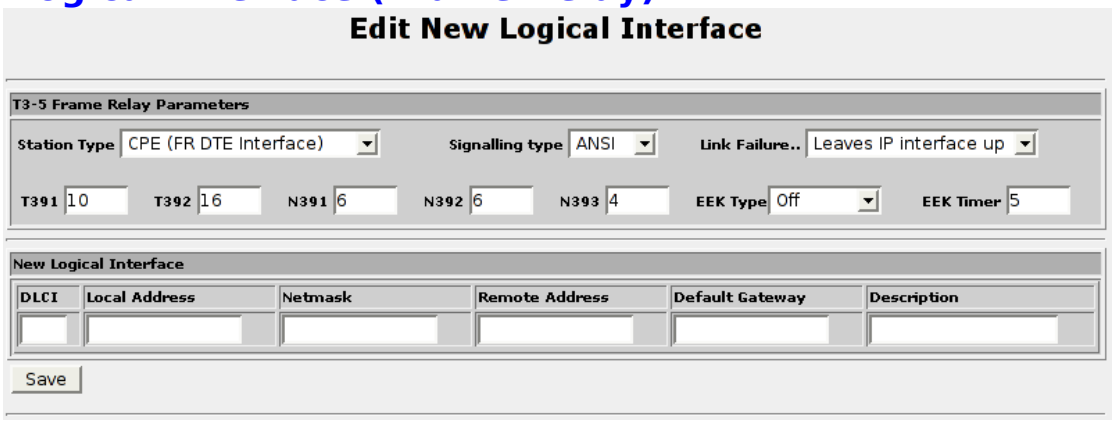

Figure 66: Edit T1 Interface

This menu allows you to display and configure logical interface fields for Frame Relay. The menu is composed of two tables. The first table provides link based configuration, which affect all DLCIs. The second table provides configuration parameters for individual DLCIs.

After the first DLCI has been configured, revisiting that DLCI will display a menu that allows additional DLCIs to be configured.

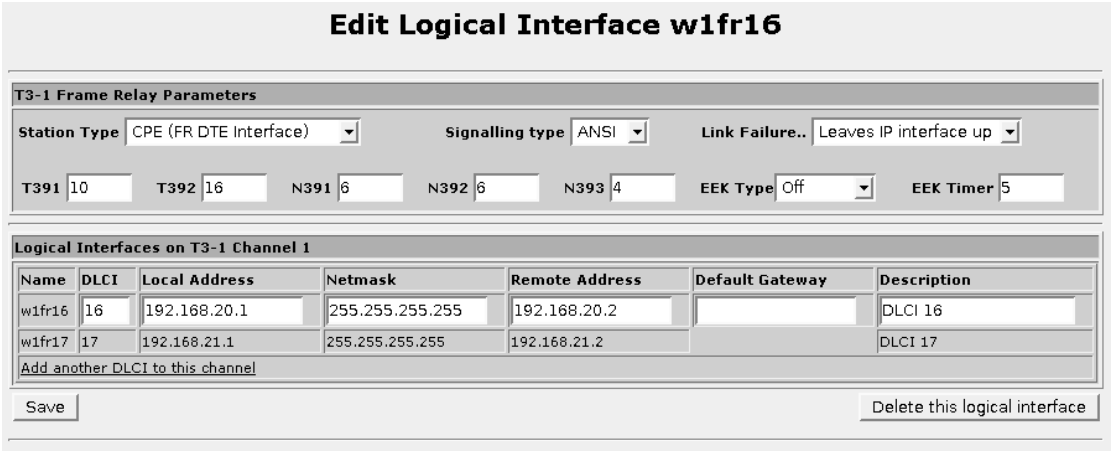

Figure 67: Edit Logical Interface (Frame Relay)

The fields and buttons in this menu are the same as those described in the **Editing A Logical Interface (Frame Relay)** section of the **Configuring Frame Relay/PPP And T1/E1** chapter.

Editing A Logical Interface (PPP)

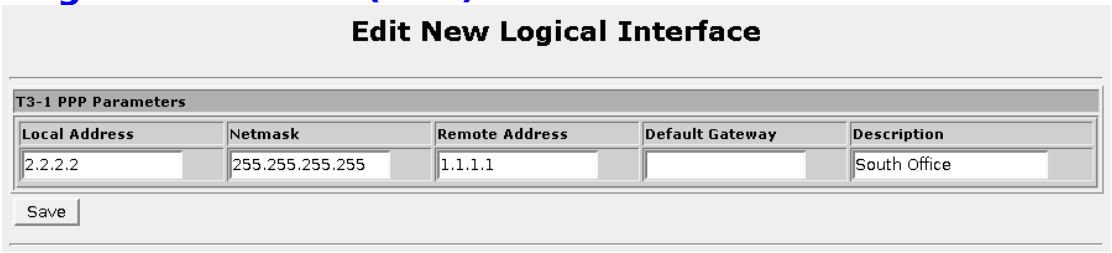

Figure 68: Edit Logical Interface (PPP)

The **Local Address**, **Netmask**, **Remote Address**, **Default Gateway and Description** fields are as described in the previous section.

T3 Statistics

When at least one logical interface is configured, T3 Link and logical interface statistics will be available. These statistics are available from links on the T3 WAN Interfaces menu.

Link Statistics are provided through the "View Link Statistics" link at the bottom of each interface table. Frame Relay and PPP statistics are available through "(Statistics)" links under the interface name column of each interface table.

Link, Frame Relay And PPP Interface Statistics are as described in the **Configuring Frame Relay/PPP And T1/E1** chapter with the exception that T3 provides only AIS, LOS, OOF and YEL alarms.

Current Routes & Interface Table

The table provided by this command is as described in the **Networking** menu, **Network Utilities** sub-menu. It is also provided here as a convenience.

Upgrading Software

For some customers, access to remote sites in accomplished solely by a T3 connection. Usually a software upgrade will stop the system being upgraded, perform the upgrade and then restart it. If T3 port was upgraded in this way, the upgrade would fail as the T3 link was taken down. Instead, T3 software upgrades modify only the software on the disk. You must schedule a reboot in order to run the new version of T3 software.

This page intentionally blank

Chapter 8 – Configuring Frame Relay/PPP And DDS

Introduction

This chapter familiarizes the user with:

- Configuring Frame Relay and PPP Links
- Viewing status and statistics
- Upgrading software

DDS Fundamentals

A Digital Data Services (DDS) line is a North American digital transmission method that operates at 56 Kbps synchronously over an unloaded, 4-Wire metallic-pair circuit.

The DDS line is typically a telephone grade network connection often called the "local loop". A Data Terminal Equipment (DTE) device attaches to the line and transmits data to the telephone company (TELCO), which routes the data to a remote DDS line. A short-haul, synchronous-data line driver known as a CSU/DSU terminates the line and attaches to the DTE. The DSU part of the DSU/CSU manages the format of the data signal while the CSU manages electrical levels, isolation and provides loopback to the TELCO.

RuggedCom DDS port provides an integrated DTE, DSU and CSU.

Location Of Interfaces And Labeling

Unlike the Ethernet ports (which are statically located), the location of T1/E1, DDS and ADSL ports in your router depends upon the number of ports and how they were ordered. Refer to the labeled hardware image as presented in the Webmin home page.

To make labeling easy to understand, all T1E1, T3, DDS and ADSL ports are assigned a unique port number that relates to the LEDs on the status panel.

LED Designations

The RuggedRouter indicates information about DDS ports on the LED Panel. A pair of LEDs will indicate traffic and link status of the port. Consult the section "Using The LED Status Panel" to determine which LEDs correspond to the port.

DDS Configuration

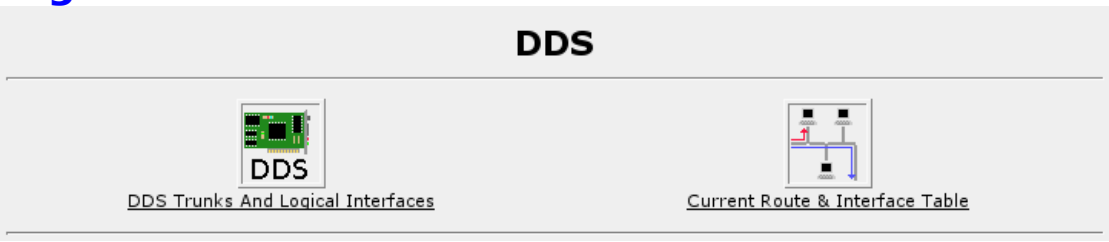

Figure 69: DDS Trunks And Interfaces

This menu allows you to display and configure DDS Trunks. The Current Routes menu will display the routes and status of the network interfaces.

DDS Network Interfaces

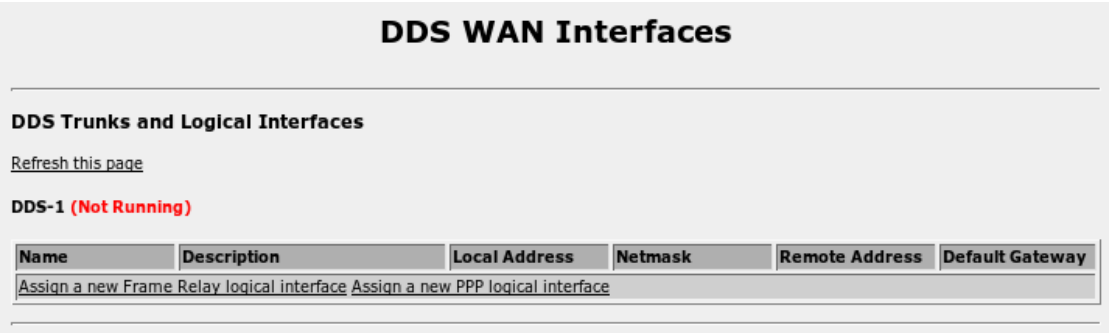

Figure 70: DDS WAN Interfaces

This menu allows you to display DDS trunks and configure the logical interfaces that run on them. A table is presented for each interface.

Interface numbers are as described by the "DDS" labels as shown in the home page chassis diagram.

The status of both the physical interface and its corresponding logical interface is shown.

If no interfaces have been configured the menu will provide links to Frame Relay and PPP configuration menus.

This menu presents connection statuses but does not update them in real time. Click on the **Refresh this page** link to update to the current status.

The menu will change after assignment of a logical interface, providing links to logical interface and link statistics.

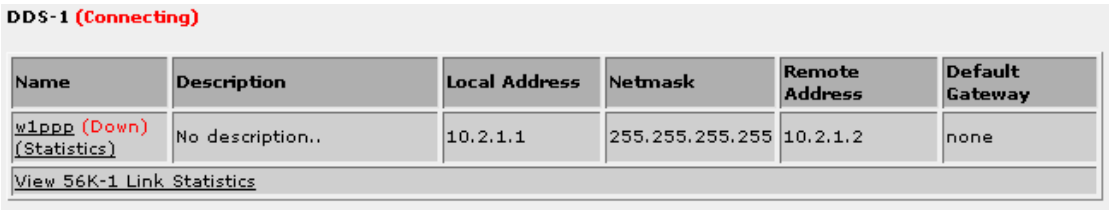

Figure 71: DDS WAN Interfaces after logical interface assignment

Naming Of Logical Interfaces

Webmin names the logical interfaces for you (but allows you to provide a description). All interfaces start with a "w" to identify them as wan interfaces, followed by the interface number. The next part of the identifier is either "ppp" or "fr" and the frame relay DLCI number.

Editing A Logical Interface (Frame Relay)

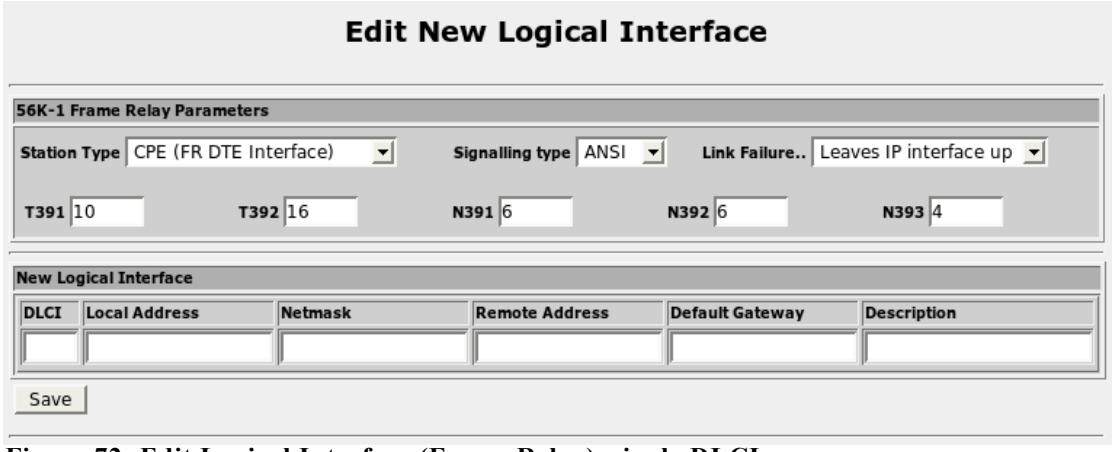

Figure 72: Edit Logical Interface (Frame Relay), single DLCI

This menu allows you to display and configure logical interface fields for Frame Relay. The menu is composed of two tables. The first table provides link based configuration, which affect all DLCIs. The second table provides configuration parameters for individual DLCIs.

After the first DLCI has been configured, revisiting that DLCI will display a menu that allows additional DLCIs to be configured.

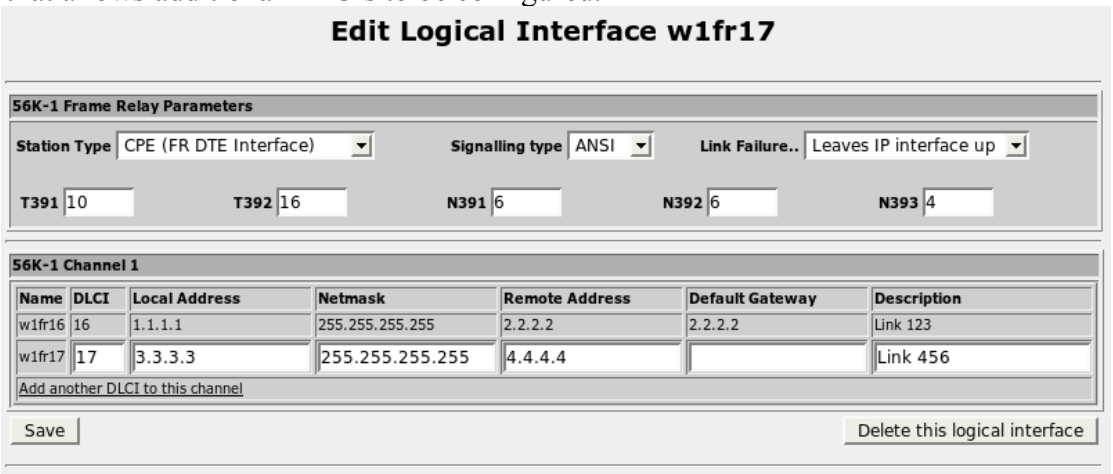

Figure 73: Edit Logical Interface (Frame Relay), multiple DLCIs

The fields and buttons in this menu are the same as those described in the **Editing A Logical Interface (Frame Relay)** section of the **Configuring Frame Relay/PPP And T1/E1** chapter.

Editing A Logical Interface (PPP)

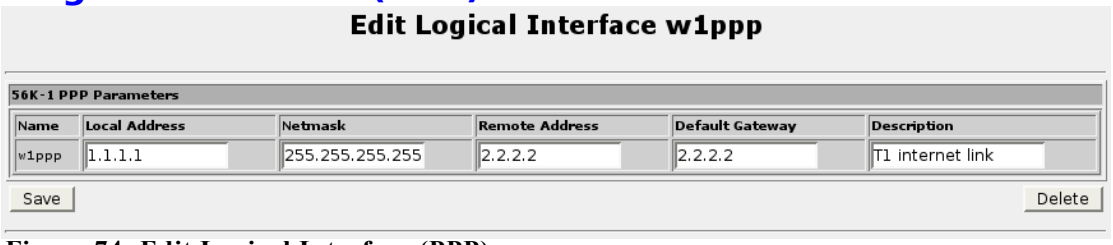

Figure 74: Edit Logical Interface (PPP)

The fields and buttons in this menu are the same as those described in the **Editing A Logical Interface (PPP)** section of the previous chapter.

DDS Statistics

When at least one logical interface is configured, DDS Link and logical interface statistics will be available. These statistics are available from links on the DDS WAN Interfaces menu.

Link Statistics are provided through the "View Link Statistics" link at the bottom of each interface table. Frame Relay and PPP statistics are available through "(Statistics)" links under the interface name column of each interface table.

Link Statistics

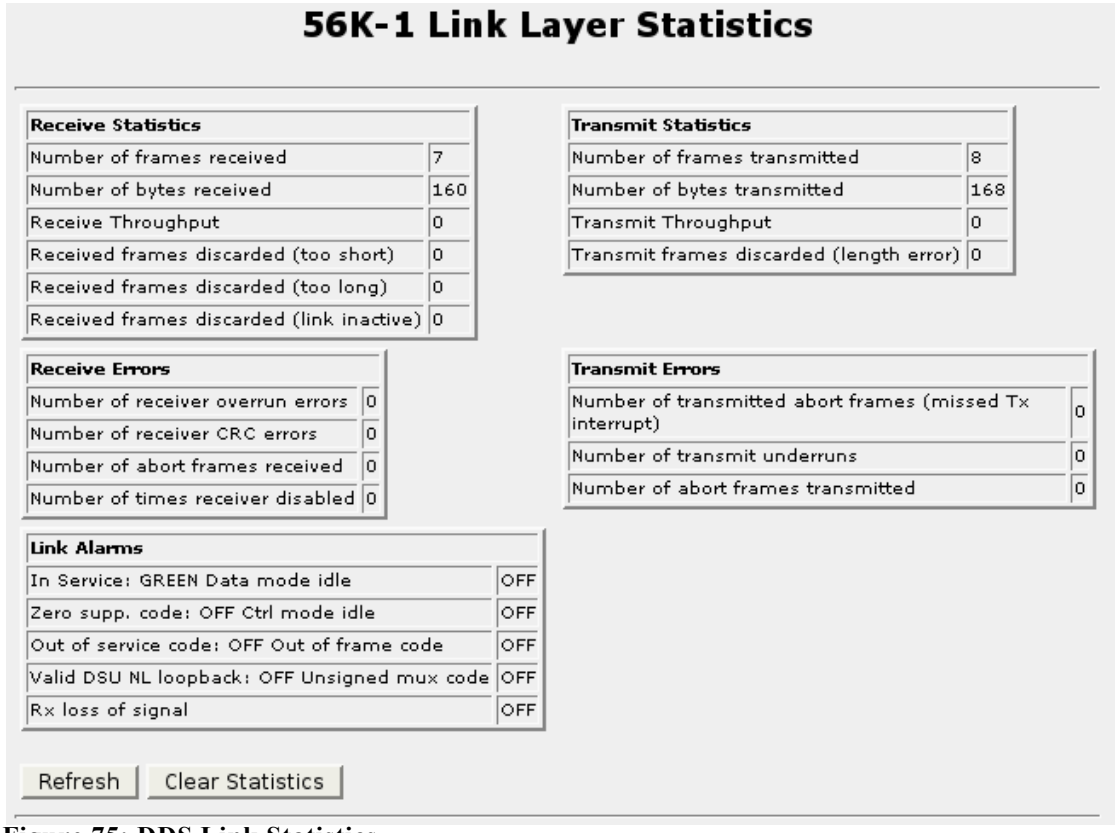

Figure 75: DDS Link Statistics

Frame Relay And PPP Interface Statistics

Frame Relay And PPP Interface Statistics are as described in the **Configuring Frame Relay/PPP And T1/E1** chapter.

DDS Loopback

When at least one logical interface is configured and that interface is active, a DDS Loopback test can be performed. This menu can be reached from a link on the DDS WAN Interfaces menu.

The remote equipment must be able to loop, allowing the entire entire line to be verified. If the remote equipment is another RuggedCom router, starting a line loopback will verify both cards and the line. DDS has no standard for performing digital loopback.

For more information on DDS loopback refer to the T1/E1 Loopback section in the chapter “Configuring Frame Relay/PPP And T1/E1”.

Current Routes & Interface Table

The table provided by this command is as described in the **Networking** menu, **Network Utilities** sub-menu. It is also provided here as a convenience.

Upgrading Software

For some customers, access to remote sites in accomplished solely by a DDS connection. Usually a software upgrade will stop the system being upgraded, perform the upgrade and then restart it. If DDS port was upgraded in this way, the upgrade would fail as the DDS link was taken down. Instead, DDS software upgrades modify only the software on the disk. You must schedule a reboot in order to run the new version of DDS software.

This page intentionally blank

Chapter 9 – Configuring PPPoE/Bridged Mode On ADSL

Introduction

This chapter familiarizes the user with:

- Configuring PPPoE and Bridged Mode Links
- Viewing status

ADSL Fundamentals

An ADSL (Asymmetric Digital Subscriber Line) line is a communications link running over regular POTS telephone service. The link is asymmetric, supporting data transfer at up to 8 Mbps from the network and up to 1 Mbps to the network. The actual bandwidth depends upon the distance between the router and telco central office, the maximum distance of which may be up to 5480 m. An ADSL card must connect to a central ADSL DSLAM for its connection.

ADSL shares ordinary telephone lines by using frequencies above the voice band. ADSL and voice frequencies will interfere with each other. If the line will be used for both data and voice, a "splitter" should be installed to divide the line for DSL and telephone.

ADSL is almost always used to make a connection to the Internet via an ISP. There are two methods for establishing the connection, PPPoe and Bridged mode.

ADSL uses the ATM protocol to communicate with the central office DSLAM. ATM uses virtual channels to route traffic and the DSL connection needs to know which virtual channels to use. Most providers use VPI=0 and VCI=35. There are exceptions to this. Some providers that use different settings are listed in the following table.

Provider	VPI	VCI
Typical Provider	0	35
Bell South	8	35
New Edge	0	38
Sprint	8	35
US West/Qwest	0	32

PPPoE/Bridged Mode Fundamentals

In PPPoE (Point-to-Point Protocol Over Ethernet) the PPP dial-up protocol is used with Ethernet over ADSL as the transport. PPPoE supports the protocol layers and authentication widely used in PPP and enables a point-to-point connection to be established in the normally multipoint architecture of Ethernet.

As your PPPoE connection is established a PPP interface will be created. The name will be "pppX" where X is the same as the interface number. Use this interface name in firewall rules.

Authentication, Addresses and DNS Servers

PPP authentication utilizes PAP or CHAP. Your ISP will provide you with a user-ID and password which you will enter in the GUI. The authentication process will assign a local IP address and addresses of the ISPs DNS servers to the router. You should use these DNS servers unless you wish to provide your own.

You will obtain either a dynamic or static IP from your ISP. Firewall configuration should be performed as is appropriate.

PPPoE MTU Issues

The use of PPPoE introduces a limitation of the maximum length of packets. The maximum Ethernet frame is 1518 bytes long. 14 bytes are consumed by the header, and 4 by the frame-check sequence, leaving 1500 bytes for the payload. For this reason, the Maximum Transmission Unit (MTU) of an Ethernet interface is usually 1500 bytes.

This is the largest IP datagram which can be transmitted over the interface without fragmentation. PPPoE adds another six bytes of overhead, and the PPP protocol field consumes two bytes, leaving 1492 bytes for the IP datagram. This reduces the MTU of PPPoE interfaces to 1492 bytes.

Packets received by hosts via Ethernet that are sized to the Ethernet MTU will be too large for the PPPoE connections MTU and will be fragmented. Large packets from hosts on the Internet will be fragmented by the ISP. The router will re-assemble these packets, but at the cost of increased latency. Configuring smaller MTUs at your hosts may reduce latency.

Bridged Mode

In bridged mode, the router simply employs the ADSL interface as a carrier of Ethernet frames. The interface will be created at boot time with a 1500 byte MTU.

No authentication information is required for bridged mode.

Your ISP will provide you with one or more IP addresses and an appropriate subnet mask. Your ISP will also suggest a DNS server which you can configure via the **Networking**, **Network Configuration**, **DNS Client** menu.

Location Of Interfaces And Labeling

Unlike the Ethernet ports (which are statically located), the location of ADSL ports in your router depends upon the number of ports and how they were ordered. Refer to the labeled hardware image as presented in the Webmin home page.

To make labeling easy to understand, all T1E1, T3, DDS and ADSL ports are assigned a unique port number that relates to the LEDs on the status panel.

LED Designations

The RuggedRouter includes two sources of LED indicated information about ADSL lines, the ADSL card itself and the LED Panel.

Four LEDs are associated with the line, next to the interface jack.

Power (Green) indicates when the card is active and powered.

Link (Green) indicates when the DSL link is established.

TX (Red) indicates when data is being transmitted over DSL.

RX (Red) indicates when data is being received over DSL.

While connecting the LEDs are flashing sequentially.

The RuggedRouter also indicates information about ADSL ports on the LED Panel. A pair of LEDs will indicate traffic and link status of the port. Consult the section "Using The LED Status Panel" to determine which LEDs correspond to the port.

ADSL Configuration

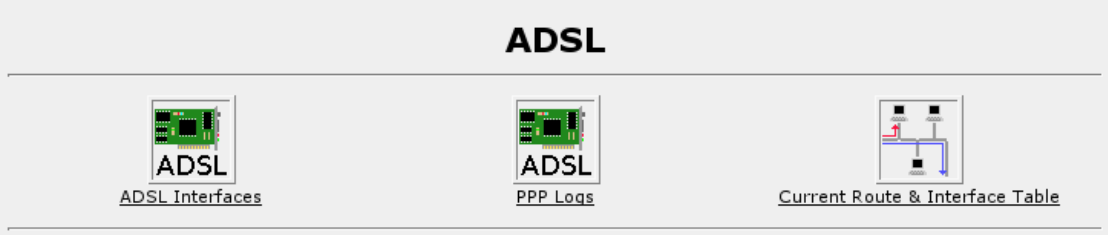

Figure 76: ADSL Interfaces

This menu allows you to display and configure ADSL interfaces. The PPP Logs menu will display a log of PPP releated information. The Current Routes menu will display the routes and status of the network interfaces.

ADSL Network Interfaces

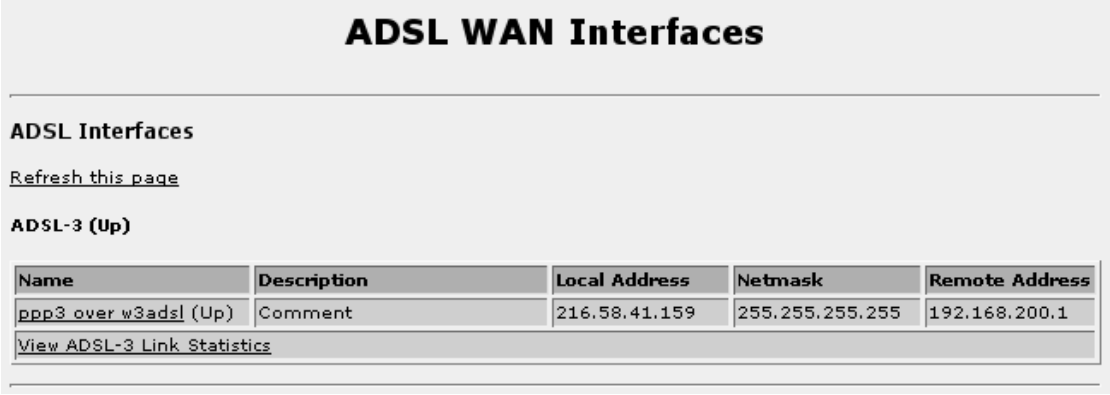

Figure 77: ADSL WAN Interfaces

This menu allows you to display and configure ADSL interfaces and the protocols that run on them. A table is presented for each interface.

Interface numbers are as described by the "ADSL" labels as shown in the home page chassis diagram.

The status of the physical interface, its corresponding logical interface and link statistics are provided.

This menu presents connection statuses but does not update them in real time. Click on the **Refresh this page** link to update to the current status.

Editing A Logical Interface (PPPoE)

Figure 78: Edit Logical Interface (PPPoE)

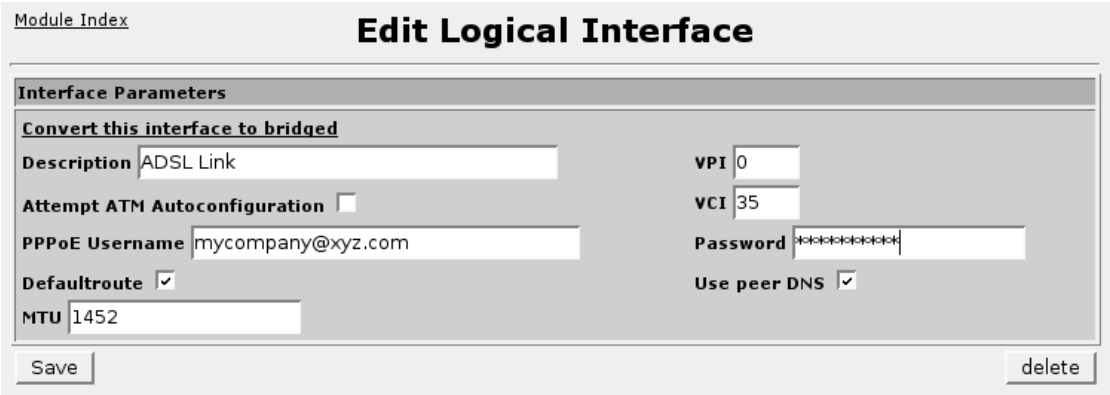

This menu allows you to display and configure logical interface fields for PPPoE and to convert the interface to Bridged Mode.

By default, interfaces are created with PPPoE. If you want the interface to be Bridged Mode, click on the **Convert this interface to bridged** link.

The **Description** field attaches a description to the logical interface viewable from the network interfaces menu.

The **VPI** field determines the VPI number the connection uses. The default of 0 is correct for most providers. The **VCI** field determines the VCI number the connection uses. The default of 35 is correct for most providers.

The **Attempt ATM Autoconfiguration** option causes the router to attempt to automatically determine the VPI and VCI used on the connection. This does not work with all providers and may case the connection to fail even if the link light is on. If this option is used it should only be used to find out what the correct values are if your provider isn't willing to help you, and when the correct values are found it should be disabled with the correct values entered in the VPI and VCI fields instead.

The **PPPoE Username** field determines the username to use when connecting to the PPPoE server as specified by your provider.

The **Password** field determines the password provided to the PPPoE server.

The **Default Route** checkbox enables automatically setting a default route using this interface whenever it connects. If this is your primary connection you probably want this option enabled.

The **Use peer DNS** checkbox enables automatically setting the DNS server entries that the PPPoE server recommends. Enable this option unless you provide your own name servers.

The **MTU** field defines the MTU size to request when connecting to the PPPoE server. In some cases the PPPoE provider may provide a smaller MTU in which case the smaller setting will be used, or it may refuse to alter the MTU and use whatever it considers to be the default.

Note: *If the negotiated MTU is different from the requested MTU, a warning will be displayed on the* ***Networking****,* ***ADSL*** *menu.*

Editing A Logical Interface (Bridged)

Figure 79: Edit Logical Interface (Bridged)

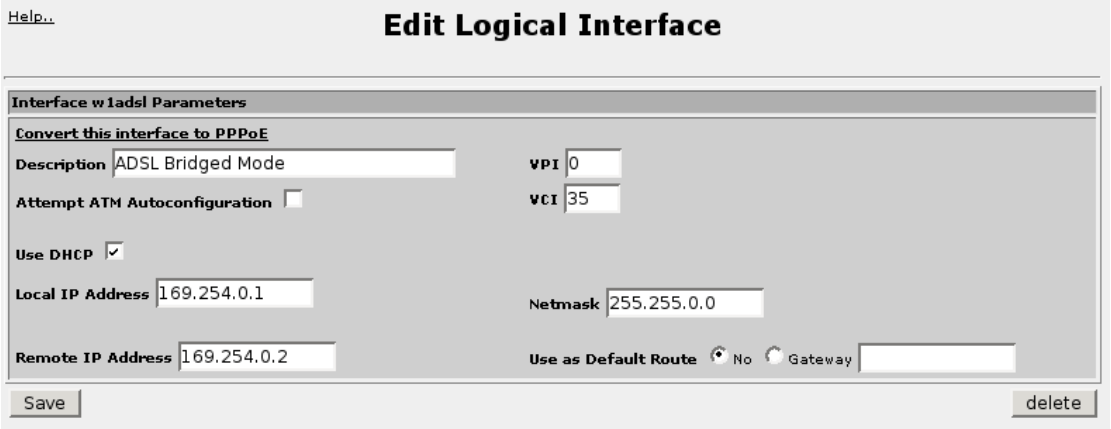

The **Description** field attaches a description to the logical interface viewable from the network interfaces menu.

The **VPI** field determines the VPI number the connection uses. The default of 0 is correct for most providers.

The **Attempt ATM Autoconfiguration** option causes the router to attempt to automatically determine the VPI and VCI used on the connection. This does not work with all providers and may case the connection to fail even if the link light is on. If this option is used it should only be used to find out what the correct values are if your provider isn't willing to help you, and when the correct values are found it should be disabled with the correct values entered in the VPI and VCI fields instead.

The **VCI** field determines the VCI number the connection uses. The default of 35 is correct for most providers.

The **Use DHCP** field forces the router to fetch its IP address from the peer via DHCP. Note that DHCP is selected the local and remote IP addresses are immediately dummied out to 169.254.0.1 and 169.254.0.2, the netmask is set to 255.255.0.0 and default gateway option is suppressed.

The **Local IP Address** field defines the IP address for this interface.

The **Netmask** field defines the network address mask The value 255.255.255.255 specifies a point-to-point connection which is almost always correct.

The **Remote IP Address** field defines the IP address for other side of this interface. As most WAN links are of point-to-point type, there is only one host connected to the other end of the link and its address is known in advance. This option is the address of the 'other end' of the link and is usually assigned by the network administrator or Internet service provider.

The **Gateway IP Address** field defines the IP address to use as the gateway for sending to other sites. This is usually the same as the Remote IP Address.

ADSL Statistics

Figure 80: ADSL Link Statistics

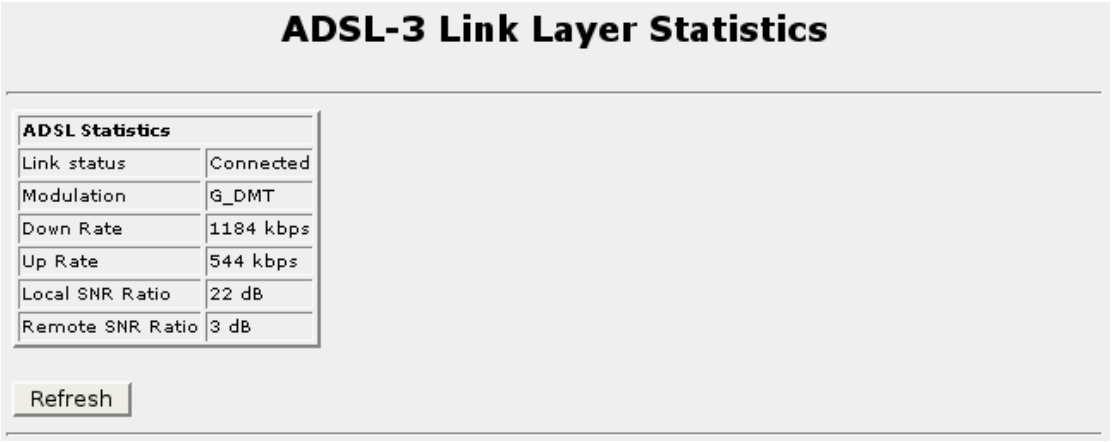

When at least one logical interface is configured, ADSL Link statistics will be available. These statistics are available from links on the DDS WAN Interfaces menu.

The **Local SNR Ratio** is an effective indicator of line quality. SNR values above 40 db correspond to excellent line quality while values below 10 db result in marginal operation or failure.

Current Routes & Interface Table

The table provided by this command is as described in the **Networking** menu, **Network Utilities** sub-menu. It is also provided here as a convenience.

Upgrading Software

For some customers, access to remote sites in accomplished solely by an ADSL connection. Usually a software upgrade will stop the system being upgraded, perform the upgrade and then restart it. If ADSL was upgraded in this way, the upgrade would fail as the ADSL link was taken down. Instead, ADSL software upgrades modify only the software on the disk. You must schedule a reboot in order to run the new version of ADSL software.

Chapter 10 – Configuring PPP and Modem

Introduction

This chapter familiarizes the user with:

- Configuring PPP Client
- Configuring PPP Server
- Configuring Dial in console
- Viewing status

Modem Fundamentals

The modem allows connections to be made over standard telephone lines. PPP is used to run network traffic over a modem link.

PPP Mode Fundamentals

PPP (Point-to-Point Protocol) is a protocol for linking two systems over a serial line.

As your PPP connection is established a PPP interface will be created. The name will be “ppp0”. Use this interface name in firewall rules.

Authentication, Addresses and DNS Servers

PPP authentication utilizes PAP or CHAP. Your ISP will provide you with a user-ID and password along with a phone number which you will enter in the GUI. The authentication process will assign a local IP address and addresses of the ISPs DNS servers to the router. You should use these DNS servers unless you wish to provide your own.

You will obtain either a dynamic or static IP from your ISP. Firewall configuration should be performed as is appropriate.

When the Modem Connects

The modem may be configured to connect at boot time.

LED Designations

The RuggedRouter provides a pair of LEDs to indicate information about the modem PPP connection.

PPP-Link will be green when the modem PPP link is established. It will flash while a connection is being established, or a console dial in session is active.

PPP-Data will flash green when there is traffic on the PPP link.

Modem Main Menu

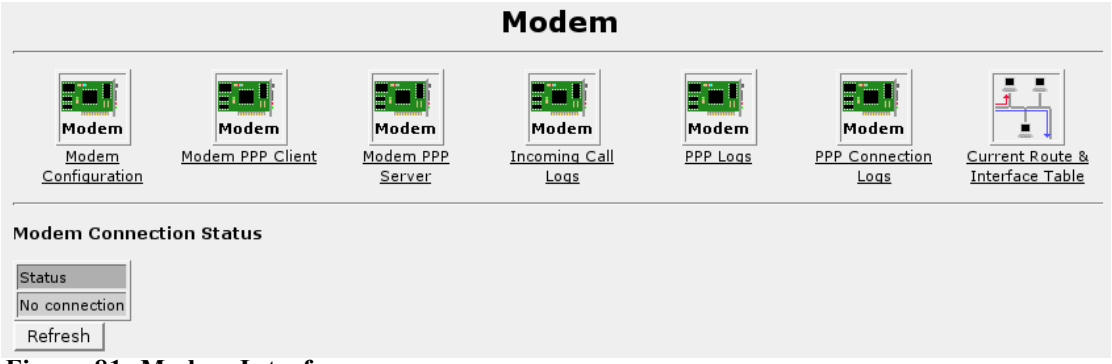

Figure 81: Modem Interface

This menu allows you to display and configure the modem interface.

Modem Configuration

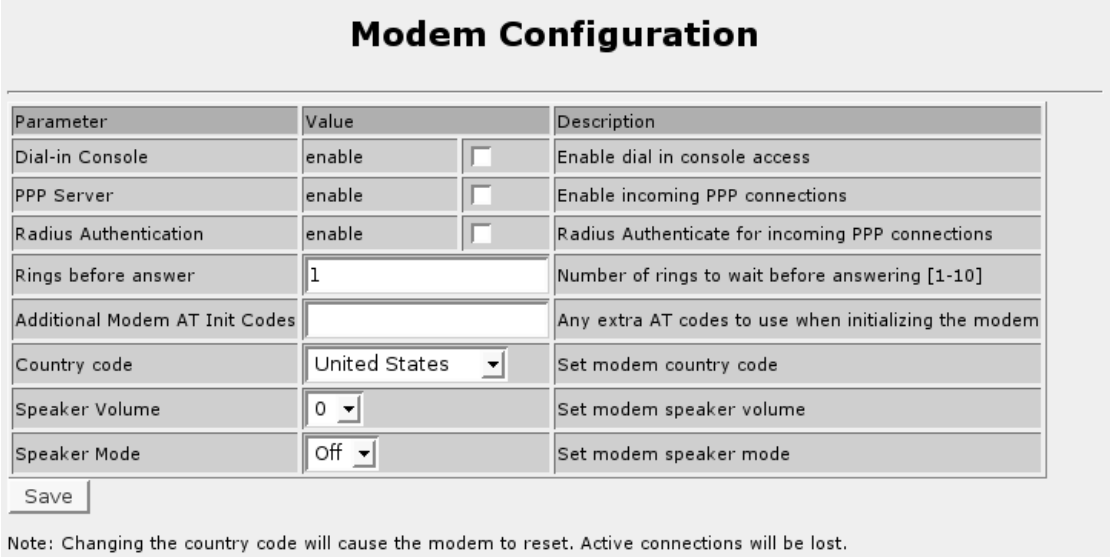

Figure 82: Edit Modem Configuration

This menu allows you to configure the modem settings and features.

The **Dial-in console** fields allows the modem to answer incoming calls and present a login just like the console serial port does. The same login is used for both.

The **PPP server** fields allows the modem to answer incoming calls and setup a PPP connection to the remote system to provide network access.

The **Radius Authentication** fields will force incoming PPP connections to authenticate against the Radius servers configured in the Maintenance menu, Radius Authentication sub-menu.

The Dial-in Console and PPP Server can be enabled at the same time. The router will automatically detect if an incoming call is PPP or console only. Is the PPP client is enabled, it will try to maintain the PPP link at all times, and hence block incoming calls most of the time. Enabling the PPP Client at the same time as the Dial-in Console and/or PPP Server is not recommended.

Rings before answer controls how many times to let the modem ring before answering the call, if Dial-in console or PPP Server is enabled.

Additional Modem AT Init Codes allows extra AT codes to be entered if required. Permitted codes are:

Blind dial

X0 - Ignore dialtone/busy signal. Blind dial.

X4 - Monitor and report dialtone/busy signal. (default)

Guard tone control

&G0 - Disable guard tone. (default)

&G1 - Enable guard tone at 550Hz.

&G2 - Enable goard tone at 1800Hz.

Pulse dialing control

&P0 - Make/break ratio of 39/61 at 10 pulses/second. (default)

&P1 - Make/break ratio of 33/67 at 10 pulses/second.

&P2 - Make/break ratio of 39/61 at 20 pulses/second.

&P3 - Make/break ratio of 33/67 at 20 pulses/second.

Compression control

%C0 - Disable data compression negotiation.

%C1 - Enable MNP5 compression negotiation.

%C2 - Enable V.42bis compression negotiation.

%C3 - Enable MNP5 and V.42bis compression negotiation. (default)

Line quality monitoring control

%E0 - Disable line quality monitor and auto-retrain.

%E1 - Enable line quality monitor and auto-retrain.

%E2 - Enable line quality monitor and fallback/fallforward. (default)

S registers

S6=X - Wait time for dialtone detection (2-255 seconds) (default=2)

S7=X - Wait time for carrier detection (1-255 seconds) (default=50)

S8=X - Pause time for comma in dial string (0-255 seconds) (default=2)

S9=X - Carrier detect response time (50-255 * .1 seconds) (default=6)

S10=X - Loss of carrier to hangup delay (50-255 * .1 seconds) (default=14)

S11=X - DTMF tone duration (50-255 * .01 seconds) (default=95)

S29=X - Hook flash dial modifier time (0-255 * .01 seconds) (default=70)

Country Code selects which country's dialing system to work with. If this is not set correctly the modem might not be able to dial or connect.

Speaker Volume controls how load the modem speaker is.

Speaker Mode controls whether the speaker on the modem is on or off.

Modem PPP Client Connections

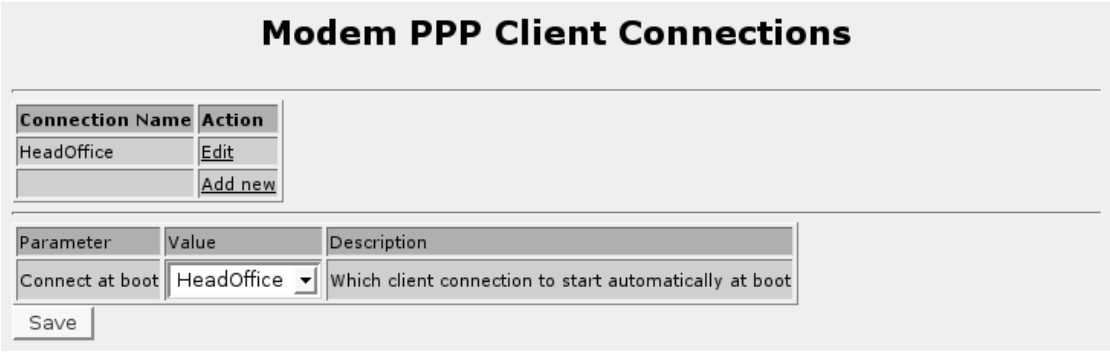

Figure 83: Configure Modem PPP Client

To edit an existing connection, click the 'Edit' link for that connection.

To create a new connection click 'Add new' link.

To have the router automatically dial a connection at boot and keep it up all the time, select which connection should be used from the drop down list of available connection profiles in the 'Connect at boot' list.

Modem PPP Client

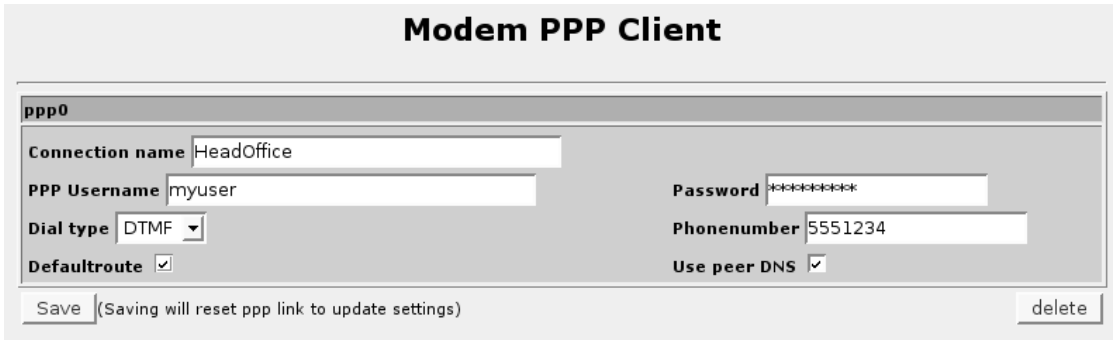

Figure 84: Configure Modem PPP Client

The **Connection Name** field determines what name will be used to refer to this connection when choosing which connection to dial automatically at boot, or which connection to use as a backup for another link.

The **PPP Username** field determines the user name to use when connecting to the PPP server as specified by your provider.

The **Password** field determines the password to use when connecting to the PPP server.

The **Dial type** field determines the type of dialing system to use on the phone line. Either DTMP (Tone dialing) or Pulse. Almost all phone systems support DTMF, and DTMF is much faster at dialing. DTMF is recommended whenever possible.

The **Phonenumber** field specifies the number to dial to connect to the PPP server.

The **Default Route** checkbox enables automatically setting a default route using this interface whenever it connects. If this is your primary connection you probably want this option enabled.

The **Use peer DNS** checkbox enables automatically setting the DNS server entries that the PPPoE server recommends. Enable this option unless you provide your own name servers.

Modem PPP Server

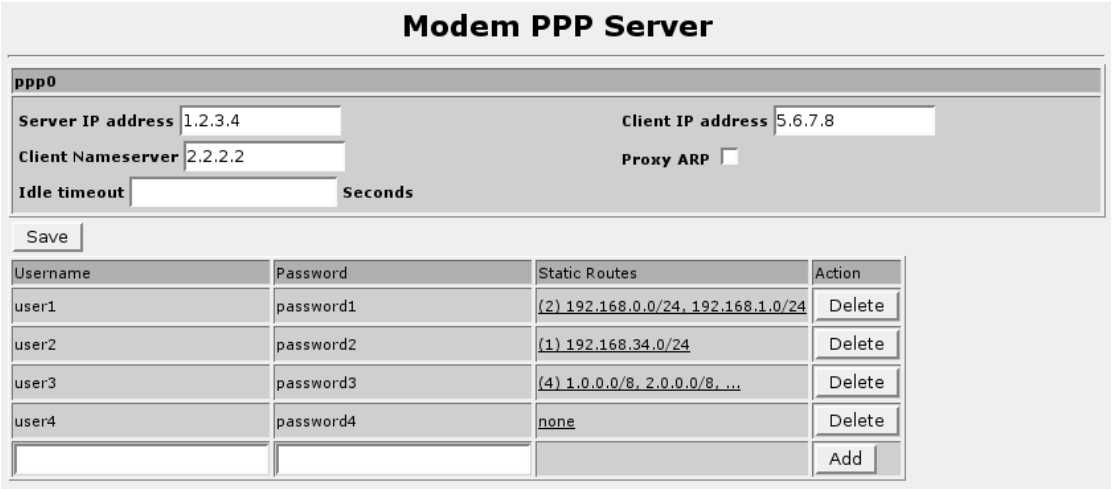

Figure 85: Configure Modem PPP Server

The **Server IP address** field controls which IP the router will use for the PPP connection.

The **Client IP address** field controls which IP to assign the to remote system which it connects.

The **Client Nameserver** field controls which nameserver (if any) the client should use for DNS lookups.

The **Proxy ARP** option makes the router attempt to proxy ARP the remote IP onto a local ethernet subnet. This requires that the Client IP address be set to an IP that would be valid on one of the ethernet subnets that the router is connected to. If this setup is used, other machines on the ethernet subnet will be able to communicate with the remote system as if it was connected directly to the ethernet subnet.

The **Idle timeout** field controls how many seconds to wait when there is no traffic on the PPP connection before hanging up the connection. Setting it to 0 or blank will disable the timeout.

The User table contains a list of users and passwords which are allowed to connect to the router by PPP. Each user can also have an optional list of subnets to create static routes to whenever their connection is established. To edit the list of routes, click on the route list for the user. To remove a user click **Delete.** To add a user, enter the username and password, and click **Add.** To change a password, enter the username and new password, and click **Add** and the password will be updated on the existing entry.

Modem Incoming Call Logs

Figure 86: Incoming Call Logs

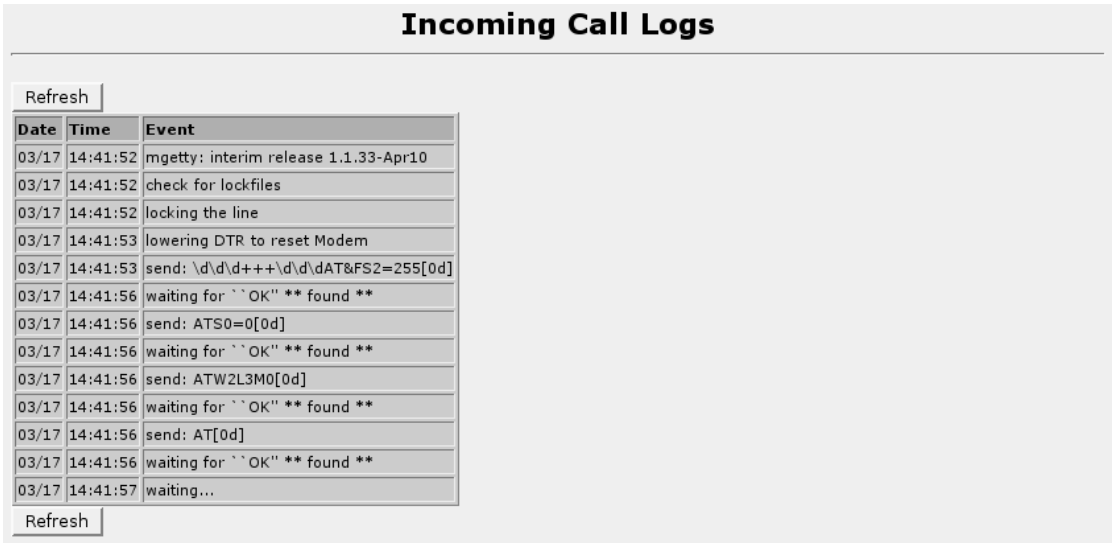

This page shows the latest log entries for incoming calls. This is mainly useful when trying to debug a problem with establishing incoming connections.

Modem PPP Logs

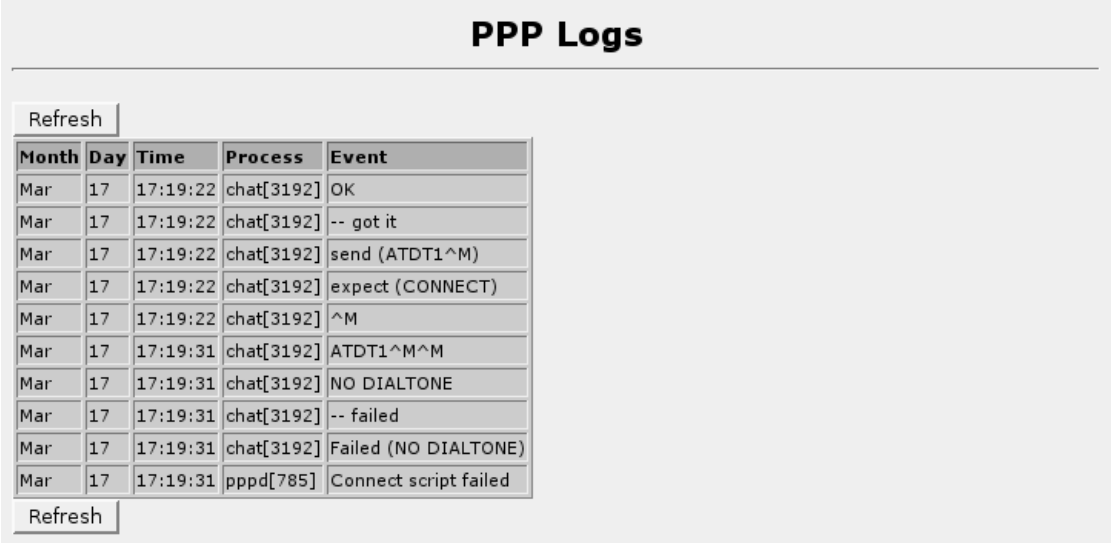

Figure 87: PPP Logs

This page shows the PPP logs. This is mainly useful when trying to debug a PPP connection problem.

Modem PPP Connection Logs

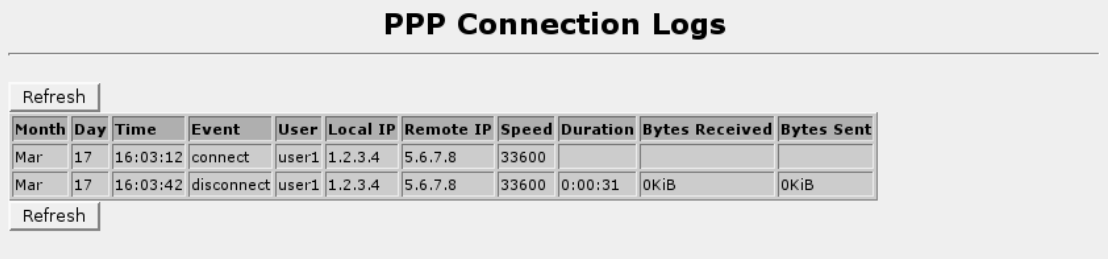

Figure 88: PPP Connection Logs

This page shows a list of PPP connections. It shows who connected, when they connected and disconnected, the connection speed, and session traffic.

Current Routes & Interface Table

The table provided by this command is as described in the **Networking** menu, **Network Utilities** sub-menu. It is also provided here as a convenience.

This page intentionally blank

Chapter 11 – Configuring The Firewall

Introduction

This chapter familiarizes the user with:

- Enabling/Disabling The Firewall
- Elements of Firewall design
- How to configure the Firewall
- Checking Firewall configuration

Firewall Fundamentals

Firewalls are software systems designed to prevent unauthorized access to or from private networks. Firewalls are most often used to prevent unauthorized Internet users from accessing private networks (intranets) connected to the Internet.

When the RuggedRouter firewall is used, the router serves a **gateway** machine through which all messages entering or leaving the intranet pass. The router examines each message and blocks those that do not meet the specified security criteria. The router also acts as a **proxy**, preventing direct communication between computers on the Internet and intranet. Proxy servers can filter the kinds of communication that are allowed between two computers and perform address translation.

Stateless vs Stateful Firewalls

Firewalls fall into into two broad categories: stateless and stateful (session-based).

Stateless or “static” firewalls make decisions about a traffic without regard to the history, simply opening a “hole” for the traffic's type (based upon TCP or UDP port number). Stateless firewalling is a relatively simple affair, easily handling web and email traffic. Stateless firewalls suffer from disadvantages, however. All holes opened in the firewall always open, there is no opening and closing connections based on outside criteria. Static IP filters offer no form of authentication.

Stateful firewalling adds considerable complexity the firewalling process by tracking the state of each connection.

A stateful firewall also looks at each packet and apply tests, but the tests applied or “rules” may be modified depending on packets that have already been processed. This is called “connection tracking”. Stateful firewalls can also recognize that traffic on connected sets of TCP/UDP ports is from a particular protocol and manage it as a whole.

Linux® netfilter, iptables And The Shoreline Firewall

The RuggedRouter employs a stateful firewall system known as **netfilter**, a set of loadable kernel modules that provides capabilities to allow session-based packet examination. The netfilter system is an interface built into the Linux kernel that allows the IP network stack to provide access to packets.

The netfilter system uses rulesets, collections of packet classification rules that determine the outcome of examination of a specific packet. The rules are defined by **iptables**, a generic table structure syntax and utility program for the configuration and control of netfilter.

In practice an iptables rule file and a script are all that are needed to load the netfilter system with rules on upon router start up. The iptables rules, however, are somewhat difficult to configure and manage.

The Shoreline Firewall, often known as shorewall, offers a more convenient approach. Shorewall is really just a front end to netfilter, maintaining the information used to generate the iptables rules in a less complicated form.

Shorewall itself does not provide a graphical front end, and instead assumes administrators will have a fair amount of familiarity with reading and editing Linux configuration files. The RuggedRouter comes with a GUI front that simplifies some of the management aspects.

Network Address Translation

Network Address Translation (NAT), enables a LAN to use one set of IP addresses for internal traffic and a second set for external traffic. The NAT function of netfilter makes all necessary IP address translations as traffic passes between the intranet and Internet. NAT is often referred to in Linux as IP Maquerading.

NAT itself provides a type of firewall by hiding internal IP addresses.

More importantly, NAT enables a network to use more internal IP addresses. Since they're used internally only, there's no possibility of conflict with IP addresses used by other organizations. Typically, your internal network will be setup to use one or more of the reserved address blocks described in RFC1918, namely:

10.0.0.0/8 (10.0.0.0 - 10.255.255.255)

172.16.0.0/12 (172.16.0.0 - 172.31.255.255)

192.168.0.0/16 (192.168.0.0 - 192.168.255.255)

As packets with these address reach the NAT gateway their source address and source TCP/UDP port number is recorded and the address/port number is translated to the public IP address and an unused port number on the public interface. When the Internet host replies to the internal machine's packets, they will be addressed to the NAT gateway's external IP at the translation port number. The NAT gateway will then search its tables and make the opposite changes it made to the outgoing packets and forward the reply packets on to the internal machine.

Translation of ICMP packets happens in a similar fashion but without the source port modification.

NAT can be used in *static* and *dynamic* modes. Static NAT masks the private IP addresses by translating each internal address to a unique external address. Dynamic NAT translates all internal addresses to one (or more) external address(es).

Port Forwarding

Port forwarding (also known as redirection) allows traffic coming from the Internet to be sent to a host behind the NAT gateway.

Previous examples have described the NAT process when connections are made from the intranet to the Internet. In those examples, addresses and ports were unambiguous.

When connections are attempted from the Internet to the intranet, the NAT gateway will have multiple hosts on the intranet that could accept the connection. It needs additional information to *identify the specific host to accept the connection*.

Suppose that two hosts, 192.168.1.10 and 192.168.1.20 are located behind a NAT gateway having a public interface of 213.18.101.62. When a connection request for http port 80 arrives at 213.18.101.62, the NAT gateway could forward the request to either of the hosts (or could accept it itself). Port forwarding configuration could be used to redirect the requests to port 80 to the first host.

Port forwarding can also remap port numbers. The second host may also need to answer http requests. As connections to port 80 are directed to the first host, another port number (such as 8080) can be dedicated to the second host. As requests arrive at the gateway for port 8080, the gateway remaps the port number to 80 and forwards the request to the second host.

Finally, port forwarding can take the source address into account. Another way to solve the above problem could be to dedicate two hosts 200.0.0.1 and 200.0.0.2 and have the NAT gateway forward requests on port 80 from 200.0.0.1 to 192.168.1.10 and from 200.0.0.2 to 192.168.1.20.

Shorewall Quick Setup

For users familiar with Shorewall the following will serves as a reminder of how to build the firewall. New users may wish to read the **ShoreWall Terminology And Concepts** section before continuing.

1) Logically partition your network into zones. Will you establish a DMZ? Will all Ethernet interfaces need to forward traffic to the public network? Which interfaces are to be treated in a similar fashion?
2) Assign your interfaces to the zones. If using T1/E1, have you created your T1/E1 interfaces prior to building the firewall?
3) Set the default policies for traffic from zone to zone to be as restrictive as possible. Has the local zone been been blocked from connecting to the DMZ or firewall? Does the DMZ or firewall need to accept connections? Which connections should be dropped and which reset? What logs are kept?
4) How is the network interface IP assigned, i.e. dynamically or statically? Do hosts at the central site need to know the local address?
5) If your network interface IP is dynamically assigned, configure masquerading.
6) If your network interface IP is statically assigned, configure Source Network address Translation (SNAT). If a sufficient number of IP addresses are provided by the ISP, static NAT can be employed instead.

7) If your hosts must accept sessions from the Internet configure the rules file to support Destination Network address Translation (DNAT). Which hosts need to accept connections, from whom and on which ports?

8) Configure the rules file to override the default policies. Have external connections been limited to approved IP address ranges. Have all but the required protocols been blocked?

9) If you are supporting a VPN, add additional rules.

10) Check the configuration using the **Shorewall Firewall** menu, "Check Firewall" button.

11) Activate the firewall. It is usually a good idea to port scan the firewall after activation and verify that logging is functioning.

ShoreWall Terminology And Concepts

This section provides background on various Shorewall terms and concepts. References are made to the section where configuration applies.

Zones

A network zone is a collection of interfaces, for which forwarding decisions are made, for example:

Name	Description
net	The Internet
loc	Your Local Network
dmz	Demilitarized Zone
fw	The firewall itself
vpn1	IPSec connections on w1ppp
vpn2	IPSec connections on w2ppp

You may create new zones if you wish. For example if all of your Ethernet interfaces are part of the local network zone, disallowing traffic from the Internet zone to the local zone will disallow it to all Ethernet interfaces. If you wanted some interfaces (but not others) to access the Internet, you could create another zone.

Zones are defined in the file /etc/shorewall/zones and are modified from the **Network Zones** menu.

Interfaces

Shorewall Interfaces are simply the Ethernet and WAN interfaces available to the router. You must place each interface into a network zone.

If an interface supports more than one subnet, place the interface in zone 'Any' and use the zone hosts setup (see below) to define a zone for each subnet on the interface.

An example follows:

Interface	Zone
eth1	loc
eth2	loc
eth3	Any
eth4	dmz
w1ppp	net

Note: *In order to improve security the router will create a zone "unusd" and unused interfaces to this zone when Shorewall starts. A policy is also installed that blocks access from "unusd" to all other zones.*

Interfaces are defined in the file /etc/shorewall/interfaces and are modified from the **Network Interfaces** menu.

Hosts

Shorewall hosts are used to assign zones to individual hosts or subnets, on an interface which handles multiple subnets. This allows the firewall to manage traffic being forwarded back out the interface it arrived on, but destined for another subnet. This is often useful for VPN setups to handle the VPN traffic separately from the other traffic on the interface which carries the VPN traffic. An example follows:

Zone	Interface	IP Address or Network
local	eth3	10.0.0.0/8
guests	eth3	192.168.0.0/24

Interfaces are defined in the file /etc/shorewall/hosts and are modified from the **Network Hosts** menu.

Policy

Shorewall policies are the default actions for connection establishment between different firewall zones. Each policy is of the form:

Source-zone *Destination-zone* *Default-action*

You can define a policy from each zone to each other. You may also use a wildcard zone of "all" to represent all zones.

The default action describes how to handle the connection request. There are six types of actions: ACCEPT, DROP, REJECT, QUEUE, CONTINUE and NONE. The first three are the most widely used and are described here.

When the **ACCEPT** policy is used, a connection is allowed. When the **DROP** policy is used, a request is simply ignored. No notification is made to the requesting client. When the **REJECT** policy is used, a request is rejected with an TCP RST or an ICMP destination-unreachable packet being returned to the client.

An example should illustrate the use of policies.

Source Zone	Destination Zone	Policy
loc	net	ACCEPT
net	all	DROP
all	all	REJECT

The above policies will:

- Allow connection requests **only** from your local network to the Internet. If you wanted to allow requests from a console on the RuggedRouter to Internet you would need to add a policy of ACCEPT fw zone to net zone.
- Drop (ignore) all connection requests from the Internet to your firewall or local network, and
- Reject all other connection requests.

Note that a client on the Internet that is probing the RuggedRouter's TCP/UDP ports will receive no responses and will not be able to detect the presence of the router. A host in the local network, on the other hand, will fail to connect to the router but will receive a notification.

Note that order of policies is important. If the last rule of this example were entered first then no connections at all would be allowed.

Policies are defined in the file /etc/shorewall/policy and are modified from the **Default Policy** menu.

Masquerading And SNAT

Masquerading and Source NAT (SNAT) are forms of dynamic NAT.

Masquerading substitutes a single IP address for an entire internal network. Use masquerading when your ISP assigns you an IP address dynamically at connection time.

SNAT substitutes a single address or range of addresses that you been assigned by your ISP. Use SNAT when your ISP assigns you one or more static IP addresses that you wish to one or more internal hosts.

The masquerading/SNAT entries are defined in the file /etc/shorewall/masq and are modified from the **Masquerading** menu. Each entry is of the form:

Interface Subnet Address Protocol Port(s)

Interface is the outgoing (WAN or Ethernet) interface and is usually your Internet interface.

Subnet is the subnet that you wish to hide. It can be an interface name (such as eth1) or an subnetted IP address.

Address is an (optional IP) address that you wish to masquerade as.

Note: *The presence of the Address field determines whether masquerading or SNAT is being used. Masquerading is used when only Interface and Subnet are present. SNAT is used when Interface, Subnet and Address are present.*

Protocol (optionally) takes on the name of protocols (e.g. tcp, udp..) that you wish to masquerade.

Ports (optionally) takes on the ports to masquerade when protocol is set to tcp or udp. These can be raw port numbers or names as found in file /etc/services.

Some examples should illustrate the use of masquerading:

Rule	Interface	Subnet	Address	Protocol	Ports
1	eth1	eth2			
2	ppp+	eth2	66.11.180.161		
3	ppp+	192.168.0.0/24	66.11.180.161		
4	w1ppp	eth1	100.1.101.16		
5	w1ppp	eth1	100.1.101.16	tcp	smtp

1) In this masquerading rule, port eth2 is connected to the local network and eth1 is connected to a DSL modem. Traffic from the subnet handled by eth2 should be translated to whatever IP is assigned to the modem. Internet clients will not be able to determine the router's public address unless some form of dynamic dns is employed.

2) In this SNAT rule a static address of 66.11.180.161 is acquired from the ISP. Traffic from the subnet handled by eth2 should be translated to 66.11.180.161 as it sent to the Internet over ppp. The + at the end of "ppp+" causes Shorewall to match any ppp interface.
3) This example is much the same as the previous one only the subnet is explicitly described, and could include traffic from any of the Ethernet ports.
4) In this SNAT rule, traffic from the subnet handled by only port eth1 should be translated to 100.1.101.16 as it sent to the Internet on t1/e1 port w1ppp.
5) This example is much the same as the previous one excepting that only smtp from eth1 will be allowed.

Masquerading and SNAT rules are defined in the file /etc/shorewall/masq and are modified from the **Masquerading** menu.

Rules

The default policies can completely configure traffic based upon zones. But the default policies cannot take into account criteria such as the type of protocol, IP source/destination addresses and the need to perform special actions such as port forwarding. The Shorewall rules can accomplish this.

The Shorewall rules provide exceptions to the default policies. In actuality, when a connection request arrives the rules file is inspected first. If no match is found then the default policy is applied. Rules are of the form:

Action Source-Zone Destination-Zone Protocol Destination-Port Source-Port Original-Destination-IP Rate-Limit User-Group

Actions are ACCEPT, DROP, REJECT, DNAT, DNAT-, REDIRECT, REDIRECT-, CONTINUE, LOG and QUEUE. The DNAT-, REDIRECT-, CONTINUE, LOG and QUEUE actions are not widely used used and are not described here.

Action	Description
ACCEPT	Allow the connection request to proceed.
DROP	The connection request is simply ignored. No notification is made to the requesting client.
REJECT	The connection request is rejected with an RST (TCP) or an ICMP destination-unreachable packet being returned to the client.
DNAT	Forward the request to another system (and optionally another port).
REDIRECT	Redirect the request to a local tcp port number on the local firewall. This is most often used to "remap" port numbers for services on the firewall itself.

The remaining fields of a rule are as described below:

Action	The action as described in the previous table.
Source-Zone	The zone the connection originated from.
Destination-Zone	The zone the connection is destined for.
Protocol	The tcp or udp protocol type.
Destination-Port	The tcp/udp port the connection is destined for.

Source-Port	The tcp/udp port the connection originated from.
Original-Destination-IP	The destination IP address in the connection request as it was received by the firewall.
Rate-Limit	A specification which allows the rate at which connections are made to be limited.
User-Group	A method of limiting outbound traffic from the firewall to a specific user, group of users and a specific application.

Some examples will illustrate the power of the rules file:

Rule	Action	Source-Zone	Destination-Zone	Protocol	Dest-Port	Source-Port	Original-Destination-IP
1	ACCEPT	net:204.18.45.0/24	fw				
2	DNAT	net	loc:192.168.1.3	tcp	ssh, http		
3	DNAT	net:204.18.45.0/24	loc:192.168.1.3	tcp	http	-	130.252.100.69
4	ACCEPT	fw	net	icmp			
5	ACCEPT	net:204.18.45.0/24	fw	icmp		8	

1) This rule accepts traffic to the firewall itself from the 204.18.45.0/24 subnet. If the default policy is to drop all requests from net to the firewall, this rule will only traffic from the authorized subnet.

2) This rule forwards all ssh and http connection requests from the Internet to local system 192.168.1.3.

3) This rule forwards http traffic from 204.18.45.0/24 (which was originally directed to the firewall at 130.252.100.69) to the host at 192.168.1.3 in the local zone. If the firewall supports another public IP address (e.g. 130.252.100.70), a similar rule could map requests to another host.

4) and 5) These rules allow the firewall to issue icmp requests to the Internet and to respond to icmp echo requests from the authorized subnet.

Rules are defined in the file /etc/shorewall/rules and are modified from the **Firewall Rules** menu.

Configuring The Firewall And VPN

Route Based Virtual Private Networking

Begin configuration by creating local, network and vpn zones (all as zone type IPV4). The openswan daemon will have created ipsecX interfaces which should be added to the vpn zone in the interfaces menu.

The IPsec protocol operates on UDP port 500 and using protocols ah (Authentication Header) and Encapsulating Security Payload (ESP) protocols. The firewall must accept this traffic in order to allow IPsec. If the firewall serves as the VPN gateway, add the following rules:

Action	Source-Zone	Destination-Zone	Protocol	Dest-Port
ACCEPT	all	fw	ah	
ACCEPT	all	fw	esp	
ACCEPT	all	fw	udp	500

IPSec traffic arriving at the firewall is directed to openswan, the IPSec daemon. Openswan then decrypts the traffic and forwards it back to shorewall on the assigned ipsecX interface. You will also need a rule to allow traffic to enter from this interface. For example, if openswan creates interface ipsec0 when its connections are established, and ipsec0 is in the zone vpn, you would need the following rule.

ACCEPT	vpn	loc

Note that if your firewall itself is required to communicate with the VPN you will need rules such as the following.

ACCEPT	vpn	fw	tcp	ssh

Policy Based Virtual Private Networking

Begin configuration by creating local, network and vpn zones. Identify the network interface that carries the encrypted IPsec traffic and make this interface part of zone "ANY" in the interfaces menu as it will be carrying both traffic for both zones.

Visit the Zone Hosts menu and, for the network interface that carries the encrypted IPsec traffic, create a zone host with zone VPN, the correct subnet and the IPsec zone option checked. If you plan to have VPN tunnels to multiple remote sites ensure that a zone host entry exists for each (or collapse them into a single subnet). Create another zone host for the same interface with a network zone, using a wider subnet mask such as 0.0.0.0/0. It is important that the vpn zone be declared before the net zone since the more specific vpn zone subnet must be inspected first.

Host Zone	Interface	Subnet	IPsec Zone?
vpn	w1ppp	192.168.1.0/24	Yes
net	w1ppp	0.0.0.0/0	No

The IPsec protocol operates on UDP port 500 and using protocols ah (Authentication Header) and Encapsulating Security Payload (ESP) protocols. The firewall must accept this traffic in order to allow IPsec.

Action	Source-Zone	Destination-Zone	Protocol	Dest-Port
ACCEPT	net	fw	ah	
ACCEPT	net	fw	esp	
ACCEPT	net	fw	udp	500

IPSec traffic arriving at the firewall is directed to openswan, the IPSec daemon. Openswan then decrypts the traffic and forwards it back to shorewall on the same interface that originally received it. You will also need a rule to allow traffic to enter from this interface.

ACCEPT	vpn	loc

Virtual Private Networking To A DMZ

If the firewall is to pass the VPN traffic through to another device (e.g. a VPN device in a DMZ) then establish a DMZ zone and install the following rules.

ACCEPT	net	dmz	ah	
ACCEPT	net	dmz	esp	
ACCEPT	net	dmz	udp	500
ACCEPT	dmz	net	ah	
ACCEPT	dmz	net	esp	
ACCEPT	dmz	net	udp	500

Firewall Main Menu

Figure 89: Starting Shorewall Firewall Menu

The above figure shows the firewall menu prior to configuration.

Configure the firewall through the provided menus. The “Check Firewall” button can be selected after each menu configuration to check the existing configuration and provide notice of items still to be configured.

When the firewall is fully configured, the “Start Firewall” button may be selected. Starting the firewall in this way will provide more detail (in the event of a problem). If the firewall starts cleanly, the menu appearance will change to that of the figure below.

In order to start the firewall at each and every boot, you must enable it via the System folder, Bootup And Shutdown menu.

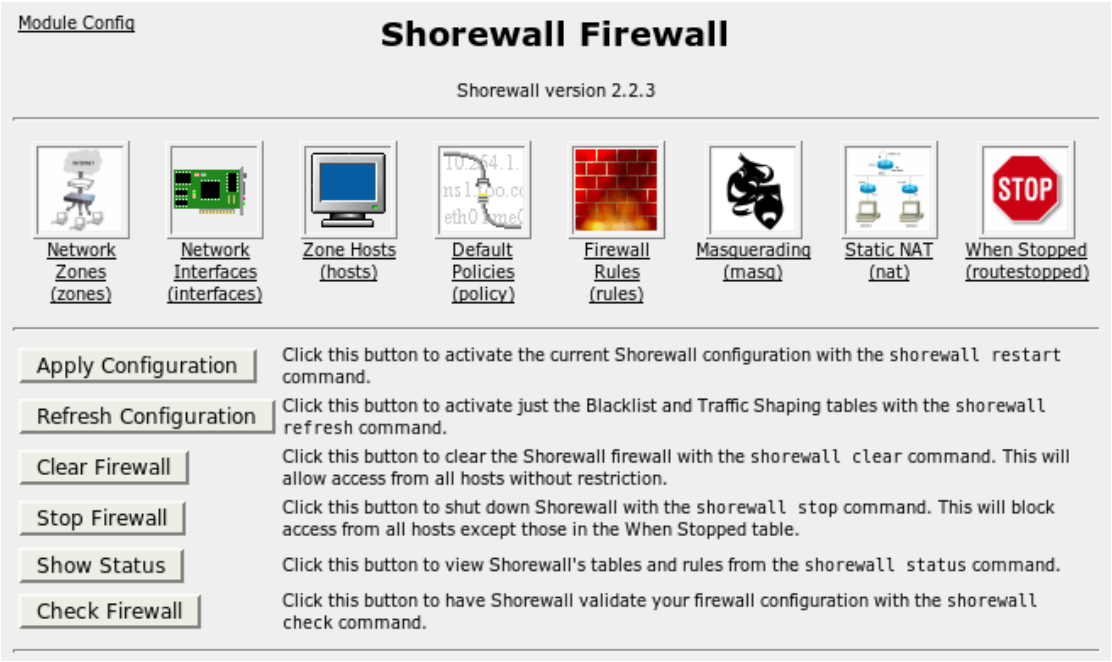

Figure 90: Shorewall Firewall Menu

The "Apply Configuration" button must be used after making configuration changes. It is recommended that the "Check Firewall" button be used first to verify that any changes made are valid.

The "Refresh Configuration" button can be used to activate changes to the blacklisted host and traffic shaping configurations.

The "Clear Configuration" button will **remove the firewall rules completely and eliminate any protection they offer**. In some cases, you might wish to do this temporarily to determine if the firewall is responsible for an application problem.

The "Stop Firewall" button will stop the firewall. **Note that you should add an entry to the "When Stopped" menu to allow access from your management station while the firewall is stopped. If you do not do this, you lose web/ssh access and have to gain access via the console in order to restart the firewall.** Stopping the firewall will not disable it. Disable the firewall via the System folder, Bootup And Shutdown menu.

The "Show Status" button presents a variety of information summarizing the status of the firewall and routing system.

The "Check Firewall" button tests the current configuration to ensure it is valid.

Network Zones

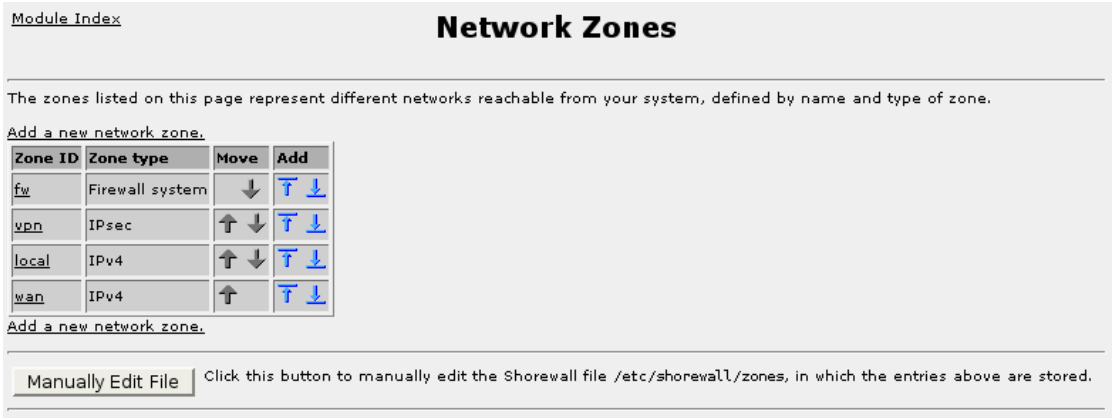

Figure 91: Firewall Network Zones

This menu allows you to add, delete and configure zones. Add a new zone by selecting the "Add a new network zone" link or by clicking on the add-above or add-below images in the **Add** field.

The **Zone Type** field controls the type of traffic carried in the zone. The Firewall system zone type is built in to the fw zone. A zone type of IPSEC is used with policy based VPNs. A zone type of IPV4 is used with normal traffic and route based VPNs.

Reorder the zones by clicking on the arrows under the **Move** field.

Note: *If you define a vpn zone whose traffic is received via a network zone, it is essential that the* vpn *zone be declared before the* network *zone.*

Clicking on a link under the **Zone ID** field will allow you to edit or delete the zone. Note that if you delete a zone you should remove any rules that reference it.

Note: There must be exactly one zone of type firewall. *Do not delete this zone.*

You may also make changes by manually editing the zone file.

Network Interfaces

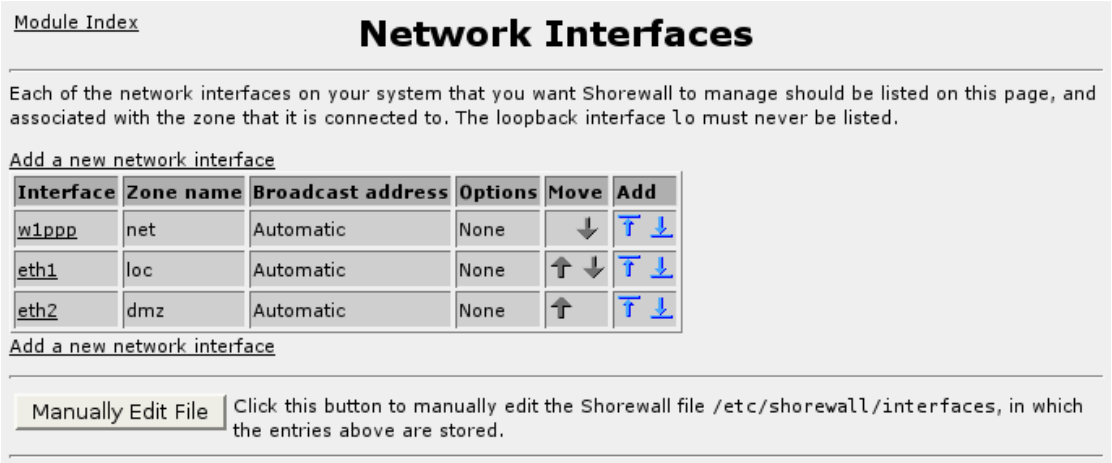

Figure 92: Firewall Network Interfaces

This menu allows you to add, delete and configure network interfaces. Add a new interface by selecting the "Add a new network interface" link or by clicking on the add-above or add-below images in the **Add** field. Reorder the interfaces by clicking on the arrows under the **Move** field.

Clicking on a link under the **Interface** field will allow you to edit or delete the interface. Note that if you delete an interface you should remove any rules that reference it.

You may also make changes by manually editing the interfaces file.

Note: *If you use a WAN interface in the firewall, the interface will be referred to by its name. Some WAN changes (such as changing the number of channels used by a T1/E1 logical interface)* ***will change the name****. Ensure that the entries in this menu reflect the correct interface names.*

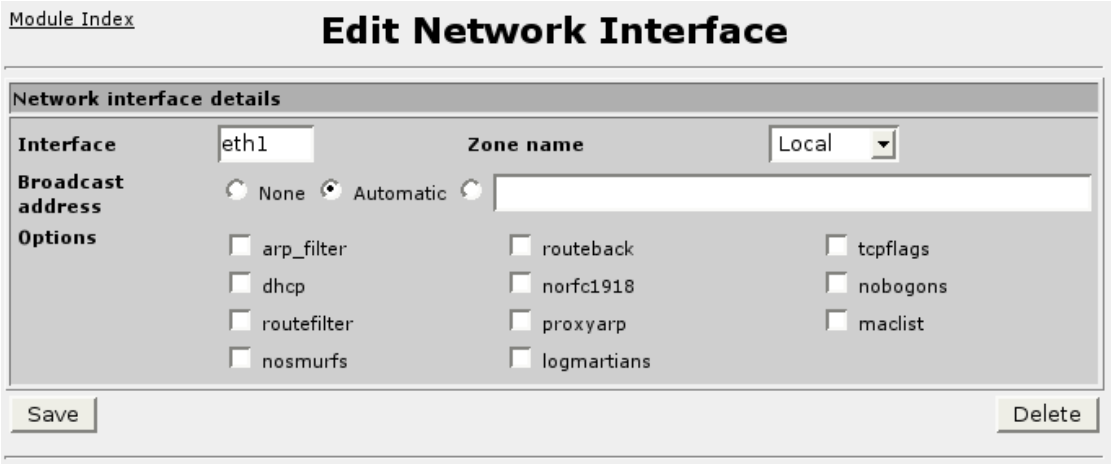

Figure 93: Editing a Firewall Network Interfaces

The **dhcp** option should be selected if interface is assigned an IP address via DHCP or is used by a DHCP server running on the firewall. The firewall will be configured to allow DHCP traffic to and from the interface even when the firewall is stopped. You may also wish to use this option if you have a static IP but you are on a LAN segment that has a lot of laptops that use DHCP and you select the **norfc1918** option (see below).

The **arp_filter** option causes this interface to only answer ARP "who-has" requests from hosts that are routed out of that interface. Setting this option facilitates testing of your firewall where multiple firewall interfaces are connected to the same HUB/Switch (all interfaces connected to the single HUB/Switch should have this option specified). Note that using such a configuration is strongly recommended against.

The **routeback** option causes Shorewall to set up handling for routing packets that arrive on this interface back out the same interface.

The **tcpflags** option causes Shorewall to make sanity checks on the header flags in TCP packets arriving on this interface. Checks include Null flags, SYN+FIN, SYN +RST and FIN+URG+PSH; these flag combinations are typically used for "silent" port scans. Packets failing these checks are logged according to the TCP_FLAGS_LOG_LEVEL option in /etc/shorewall/shorewall.conf and are disposed of according to the TCP_FLAGS_DISPOSITION option.

The **norfc1918** option causes packets arriving on this interface and that have a source or destination address that is reserved in RFC 1918 to be dropped after being optionally logged.

The **nobogons** option causes packets arriving on this interface that have a source address reserved by the IANA or by other RFCs (other than 1918) to be dropped after being optionally logged.

The **routefilter** option invokes the Kernel's route filtering (anti-spoofing) facility on this interface. The kernel will reject any packets incoming on this interface that have a source address that would be routed outbound through another interface on the firewall.

The **proxyarp** option causes Shorewall to set proxy arp for the interface. Do **not** set this option if implementing Proxy ARP through entries in /etc/shorewall/proxarp.

The **maclist** option causes all connection requests received on this interface to be subject to MAC address verification. May only be specified for Ethernet interfaces.

The **nosmurfs** option causes incoming connection requests to be checked to ensure that they do not have a broadcast or multicast address as their source. Any such packets will be dropped after being optionally logged according to the setting of SMURF_LOG_LEVEL in /etc/shorewall/shorewall.conf.

The **logmartians** option causes the martian logging facility will be enabled on this interface. See also the LOG_MARTIANS option in /etc/shorewall/shorewall.conf.

Network Zone Hosts

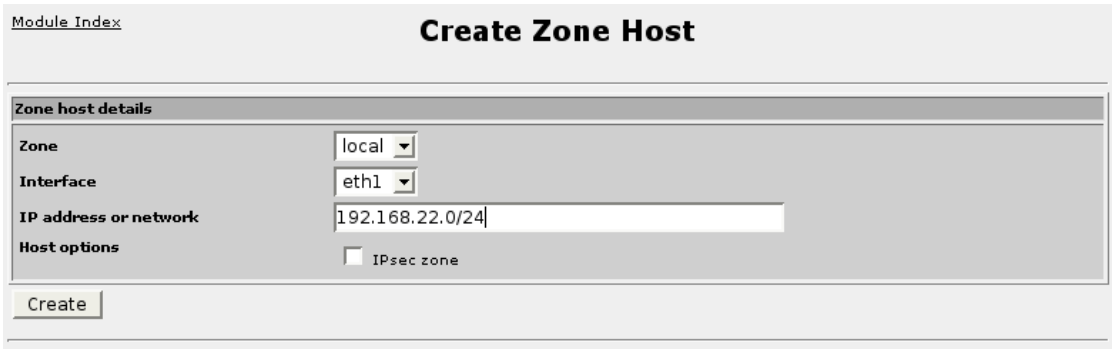

Figure 94: Firewall Zone Hosts

This menu allows you to add, delete and configure interfaces hosting multiple zones. Add a new zone host by selecting the “Add a new zone host” link or by clicking on the add-above or add-below images in the **Add** field. Reorder the hosts by clicking on the arrows under the **Move** field.

The **Zone** field selects a zone that will correspond to a subnet on the interface in question. The **Interface** field describes that interface and the **IP address or network** field describes the subnet.

Selecting the **IPSEC zone Host Option** field will identify that the traffic to host in this zone is encrypted.

The **Save** and **Delete** buttons will allow you to edit or delete the zone host. You may also make changes by manually editing the policy

Default Policies

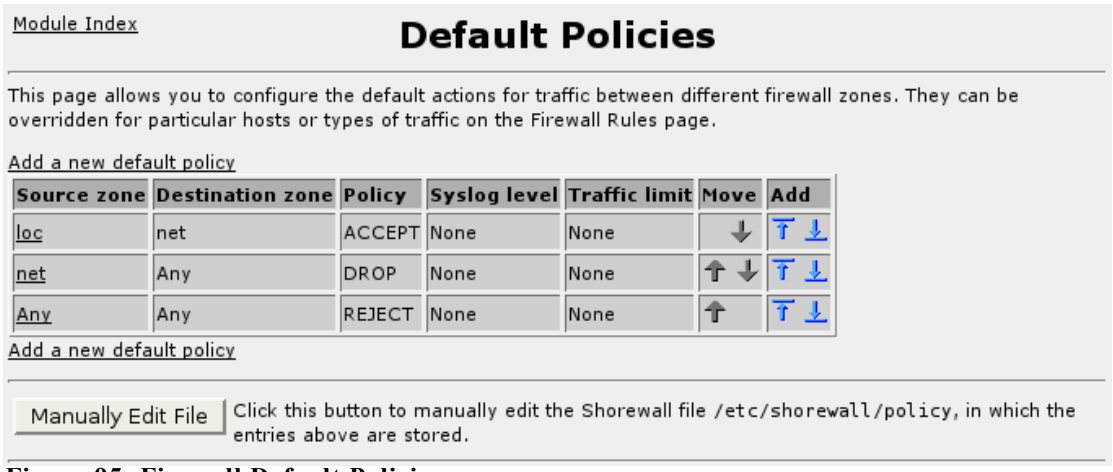

Figure 95: Firewall Default Policies

This menu allows you to add, delete and configure default policies. Add a new policy by selecting the "Add a new default policy" link or by clicking on the add-above or add-below images in the **Add** field. Reorder the policies by clicking on the arrows under the **Move** field.

Clicking on a link under the **Source zone** field will allow you to edit or delete the policy, as shown below. You may also make changes by manually editing the policy file.

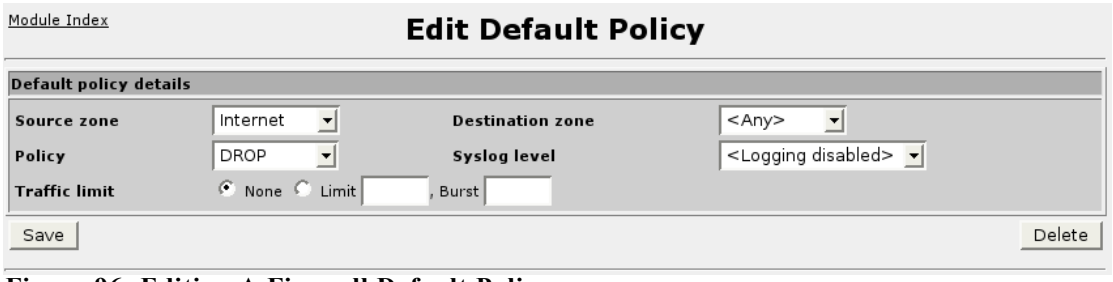

Figure 96: Editing A Firewall Default Policy

The **Syslog level** field causes a log entry to be generated every time the rule is followed.

The **Traffic limit** fields allow you to place an upper limit upon the rate at which the rule is applied. The **Limit** field is the steady state rate and is of the form "X/sec" or "X/min" where X is the number of allowed rule followings. The **Burst** field denotes the largest permissible burst and defaults to five if not configured.

Masquerading

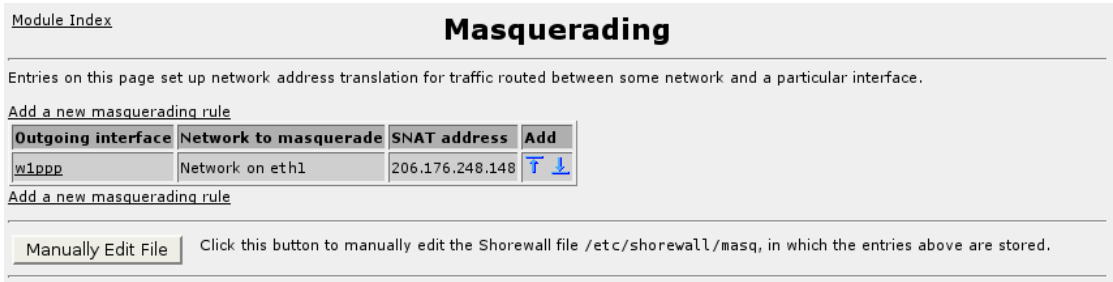

Figure 97: Firewall Masquerading And SNAT

This menu allows you to add, delete and configure masquerading and SNAT rules. Add a new rule by selecting the "Add a new masquerading rule" link or by clicking on the add-above or add-below images in the **Add** field. Reorder the policies by clicking on the arrows under the **Move** field.

Clicking on a link under the **Outgoing interface** field will allow you to edit or delete the rule, as shown below. You may also make changes by manually editing the rule file.

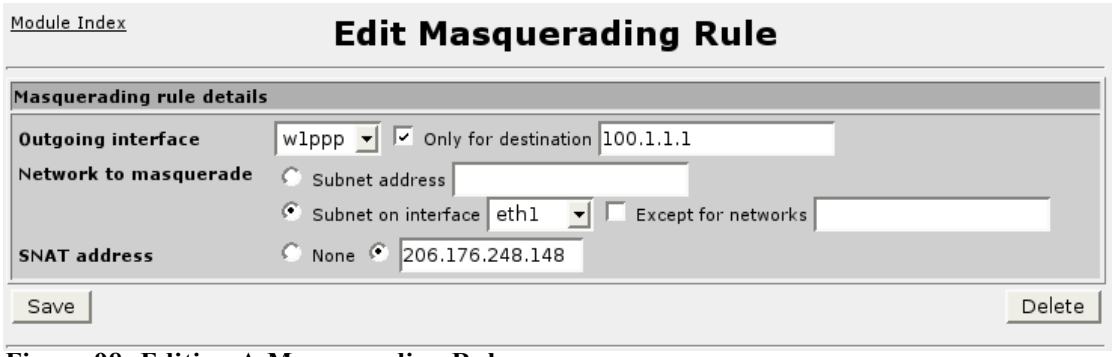

Figure 98: Editing A Masquerading Rule

The **Only for destination** field restricts the masquerading to the specified IP address.

The **Network to masquerade** fields determine the interface or subnet on the private network that you wish to masquerade. The **Except for networks** field restricts traffic from the specified subnet.

The **SNAT address** field is used to determine whether masquerading or SNAT is being performed. If checked, the entered IP address is used as a SNAT address.

Firewall Rules

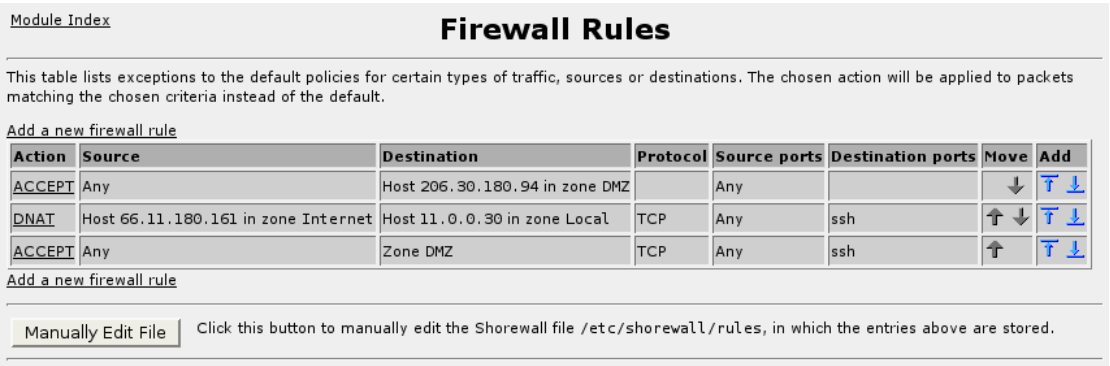

Figure 99: Firewall Rules

This menu allows you to add, delete and configure firewall rules. These rules are inspected and applied before the default policies are used. Add a new rule by selecting the "Add a new firewall rule" link or by clicking on the add-above or add-below images in the **Add** field. Reorder the policies by clicking on the arrows under the **Move** field.

Clicking on a link under the **Action** field will allow you to edit or delete the rule, as shown below. You may also make changes by manually editing the rule file.

Figure 100: Editing A Firewall Rule

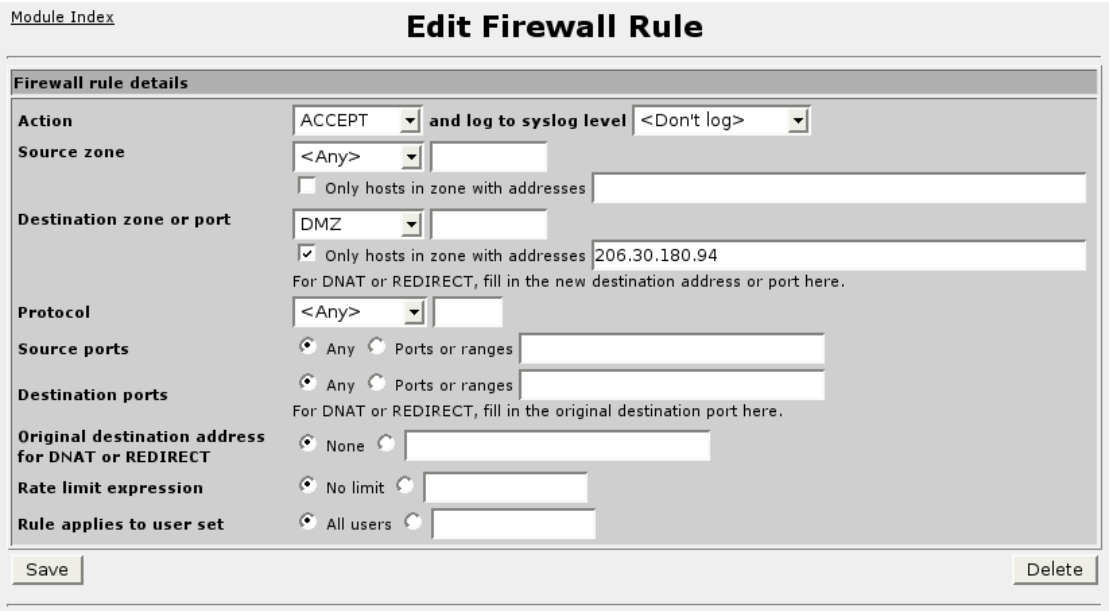

The following fields describe the information to match against the incoming connection request in order to apply this rule.

The **Action** field specifies the final action of the rule. The **and log to syslog** field determines whether logging will take place and at which logging level.

The **Source zone** field specifies the zone the request originates from.

The **Destination zone or port** field specifies the requests destination zone.

The **Protocol** field specifies the protocol (tcp, udp or icmp) to match.

The **Source ports** and **Destination ports** fields specifies the requests tcp or udp port numbers to match.

The **Original destination address** field matches the requests destination IP address.

Note: *If you use are using DNAT to port forward, enter the original destination address here and the forwarded address in the* ***Destination zone or port*** *fields* ***Only hosts in zone with address*** *sub-field.*

The **Rate limit expression** fields specifies a rate limit control of the form "X/sec" or "X/min" where X is the number of allowed requests in the time period. A burst limit field ":Y"where Y is the maximum consecutive number of requests and defaults to five if not configured.

The **Rule applies to user set** fields allow advanced users to match the rule against specific users and groups. This matching only takes place when the source of the traffic is the firewall itself.

Static NAT

Figure 101: Static NAT

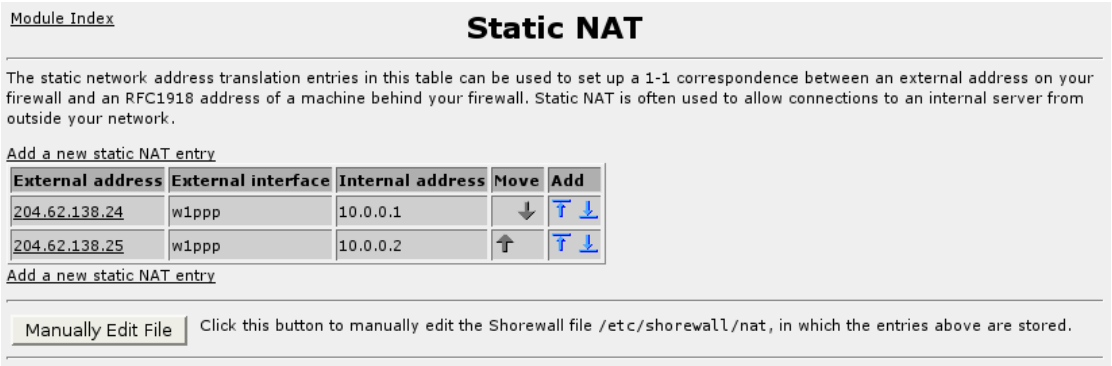

This menu allows you to add, delete and static NAT translations. Add a new translations by selecting the "Add a new static NAT entry" link or by clicking on the add-above or add-below images in the **Add** field. Reorder the translations by clicking on the arrows under the **Move** field.

Clicking on a link under the **External Address** field will allow you to edit or delete the rule, as shown below. You may also make changes by manually editing the rule file.

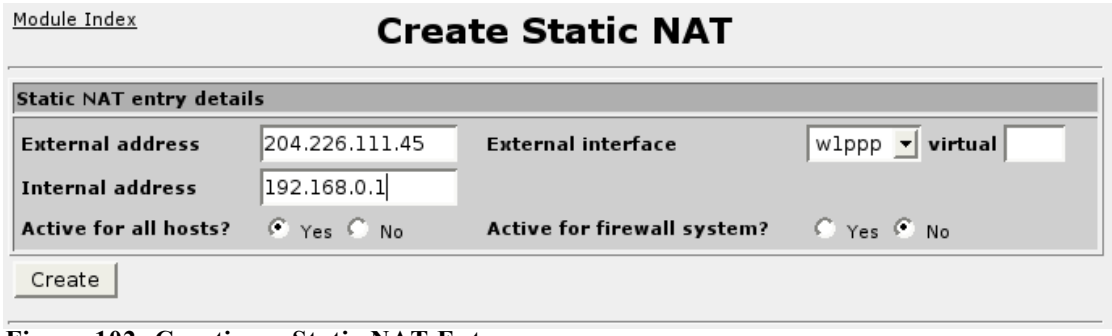

Figure 102: Creating a Static NAT Entry

The **External address** and **Internal address** fields specify the addresses to translate. describe the information to match against the incoming connection request in order to apply this rule.

The **External interface** field specifies the interface to perform the translation upon.

The **Active for all hosts** field is used to specify whether access to the external IP from all firewall interfaces should undergo NAT (Yes or yes) or if only access from the interface in the INTERFACE column should undergo NAT.

The **Active for firewall system** field is used to specify whether packets originating from the firewall itself and destined for the EXTERNAL address are redirected to the internal ADDRESS.

Actions When Stopped

Figure 103: Actions When Stopped

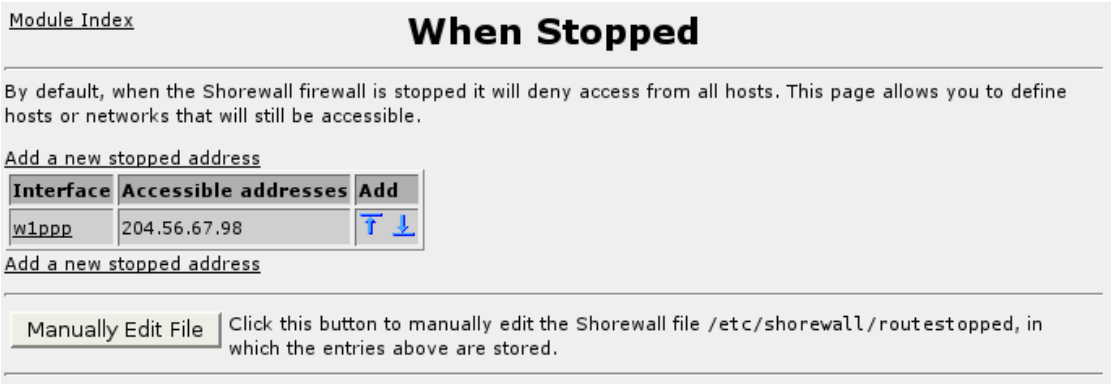

This menu allows you to control which addresses the firewall will accept connections from after it has been stopped. Add a new translations by selecting the "Add a new stopped address" link or by clicking on the add-above or add-below images in the **Add** field. Reorder the translations by clicking on the arrows under the **Move** field.

Clicking on a link under the **Interface** field will allow you to edit or delete the rule, as shown below. You may also make changes by manually editing the rule file.

This page intentionally blank

Chapter 12 – Configuring An IPsec VPN

Introduction

This chapter familiarizes the user with:

- Configuring IPsec VPN Global Options
- Creating VPN Connections
- Enabling And Starting IPsec
- Obtaining VPN Status

VPN Fundamentals

IPsec (Internet Protocol SECurity) uses strong cryptography to provide both authentication and encryption services. Authentication ensures that packets are from the right sender and have not been altered in transit. Encryption prevents unauthorized reading of packet contents.

These services allow you to build **secure tunnels through untrusted networks**. Everything passing through the untrusted network is encrypted by the IPsec gateway and decrypted by the gateway at the other end. The result is a Virtual Private Network (VPN), a network which is effectively private even though it includes machines at several different sites connected by the insecure Internet.

The IPsec protocols were developed by the Internet Engineering Task Force (IETF) and will be required as part of IPv6, the next generation.

Openswan is the open source implementation of IPsec used by the RuggedRouter.

The protocols used by IPsec are the Encapsulating Security Payload (ESP) and Internet Key Exchange (IKE) protocols.

ESP provides encryption and authentication (ensuring that a message originated from the expected sender and has not been altered on route).

IKE negotiates connection parameters, including keys, for ESP. IKE is based on the Diffie-Hellman key exchange protocol, which allows two parties without any initial shared secret to create one in a manner immune to eavesdropping.

IPsec Modes

IPSec has two basic modes of operation. In *transport mode*, IPSec headers are added as the original IP datagram is created. The resultant packet is composed of an IP header, IPSec headers and IP payload (including a transport header). Transport mode is most commonly used between IPsec end-stations, or between an end-station and a gateway/

In *tunnel mode*, the original IP datagram is created normally and then encapsulated into a new IP datagram. The resultant packet is composed of an new IP header, IPSec headers, old IP header and IP payload. Tunnel mode is most commonly used between gateways, the gateway acting as a proxy for the hosts behind it.

Policy Vs Route Based VPNs

The RuggedRouter supports two main modes of VPN: policy and route based VPN.

With route based VPNs:

- Openswan generates an IPSEC interface for each VPN tunnel,
- As the tunnel is brought up a route for the subnet at the other end of the tunnel is created through that interface,
- Any traffic destined for tunnel's remote subnet is forwarded to the IPSEC interface and encoded and transmitted,
- The firewall is configured with a vpn zone (zone type IPV4), the IPSEC interface is included in the zone,
- As IPsec packets are received, openswan decodes them and directs the decoded packet to the IPSEC interface,
- Firewalling can be performed by simply accepting all traffic to and from the zone containing the IPSEC interfaces,
- It is possible to use a tunnel to provide the default route by making the subnet at the other end of the tunnel be 0.0.0.0/0.

With policy based VPNs:

- Openswan will not generate IPSEC interfaces,
- The routing table is not involved in deciding which packets should go to the ipsec layer,
- Only traffic matching the tunnel's local and remote subnets are forwarded to it. Normal traffic is routed by one set of rules and VPN traffic is routed based on different rules,
- The firewall is configured with a vpn zone of zone type IPSEC,
- As IPsec packets are received, openswan decodes them, policy flags them as IPSEC encoded and presents them as arriving on the same interface they originally arrived at.
- Firewall rules must be written to allow traffic to and from tunnels based upon the the normal form of source/destination IP addresses and IP protocol and port numbers. These, by virtue of the zones they match, use the policy flagging inserted by netkey and routes them to the proper interface.

Route based VPNs are the default. This type of VPN is recommended as it is simpler to configure.

Supported Encryption Protocols

Openswan supports the following standard encryption protocols:

- 3DES (Triple DES) – Uses three DES encryptions on a single data block, with at least two different keys, to get higher security than is available from a single DES pass. 3DES is the most CPU intensive cipher.
- AES – The Advanced Encryption Standard protocol cipher uses a 128-bit block and 128, 192 or 256-bit keys. This is the most secure protocol in use today, and is much preferred to 3DES due to its efficiency.

Public Key And Pre-shared Keys

In **public key cryptography**, keys are created in matched pairs (called public and private keys). The public key is made public while the private key is kept secret. Messages can then be sent by anyone who knows the public key to the holder of the private key. Only the owner of the private key can decrypt the message.

When you want to use this form of encryption, each router configures its VPN connection to use the RSA algorithm and includes the public signature of its peer. The RuggedRouter's public signature is available from the output of the **Show Public Keys** menu.

In **secret key cryptography**, a single key known to both parties is used for both encryption and decryption.

When you want to use this form of encryption, each router configures its VPN connection to use a secret pre-shared key. The pre-shared key is configured through the **Pre-shared Keys** menu.

Note: *Use of pre-shared keys require that the IP addresses of both ends of the VPN connection be statically known, so they can't be used with sites with dynamic IPs.*

X509 Certificates

When one side of the VPN connection is placed from a dynamic IP (the so-called "roaming client"), X509 Certificates may be used to authenticate the connection. Certificates are digital signatures that are produced by a trusted source, namely a Certificate Authority (CA). For each host, the CA creates an certificate that contains CA and host information and "signs" the certificate by creating a digest of all the fields in the certificate and encrypting the hash value with its **private key**. The encrypted digest is called a "digital signature". The host's certificate and the CA **public key** are installed on all gateways that the host connects to.

When the gateway receives a connection request it uses the CA **public key** to decrypt the signature back into the digest. It then recomputes its own digest from the plain text in the certificate and compares the two. If both digests match, the integrity of the certificate is verified (it was not tampered with), and the public key in the certificate is assumed to be the valid public key of the connecting host.

NAT Traversal

Historically, IPSec has presented problems when connections must traverse a firewall providing Network Address Translation (NAT). The Internet Key Exchange (IKE) used in IPSec is not NAT-translatable. When IPSec connections must traverse a firewall IKE messages and IPSec-protected packets must be encapsulated as User Datagram Protocol (UDP) messages. The encapsulation allows the original untranslated packet to be examined by IPSec.

Other Configuration Supporting IPSec

If the router is to support a remote IPSec client and the client will be assigned an address in a subnet of a local interface, you must activate proxy ARP for that interface. This will cause the router to respond to ARP requests on behalf of the client and direct traffic to it over its connection.

IPSec relies upon the following protocols and ports:

- protocol 51, IPSEC-AH Authentication Header (RFC2402),
- protocol 50, IPSEC-ESP Encapsulating Security Payload (RFC2046),
- UDP port 500.

You must configure the firewall to accept connections on these ports and protocols. See the **Configuring The Firewall** chapter, **Configuring The Firewall And VPN** section for details.

The Openswan Configuration Process

Each VPN connection has two ends, in the local router and the remote router. The Openswan developers designed the configuration in such a way that the configuration record describing a VPN connection can be used without change at either end. One side of the connection (typically the local side) is designated the "left" side and the other is designated the "right" side.

A convenient method is to configure both ends simultaneously, having two browser windows up. The relevant information is cut and pasted from window to window.

This module also includes tools to export and import the connection data. The configuration can thus be generated at one router, exported, and imported at the remote router.

IPsec and Router Interfaces

The IPsec daemon requires router interfaces to exist before it starts. If none of the interfaces needed by IPsec exist, IPsec will check for them every minute until at least one does.

Note that in the unlikely event that IPsec uses multiple network interfaces, a stop of any of those interfaces will cause all tunnels to stop.

IPsec may have to be manually restarted after configuring network interfaces when multiple tunnels exist.

VPN Main Menu Before Key Generation

Figure 104: IPsec VPN Configuration Menu Before Key Generation

Upon the first entry to this menu you will prompted to generate a VPN host key. Key generation will require about 30 seconds to complete after which the menu appearance will change.

VPN Main Menu

The new menu appearance will resemble that of the following menu with the exception that you will be warned that VPN networking is not enabled. Enable VPN networking via the System folder, Bootup And Shutdown menu.

Figure 105: IPsec VPN Configuration Menu Before After Generation

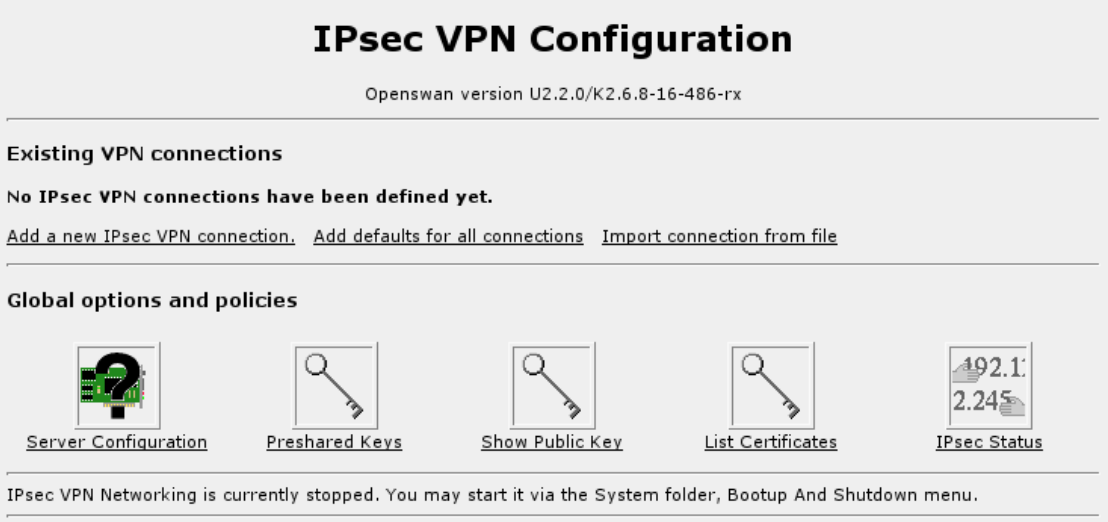

After a VPN connection is created this menu will display an icon for the connection, as shown in the next view of the VPN Configuration menu.

The “Add defaults for all connections” link allows you to create a profile that will apply to all connections for items such as key type, encryption protocol and compression. These defaults can then be overridden on a per connection basis.

The “Add a new IPsec VPN connection” link creates a new connection and its icon.

The “Import connection from file” link creates new connections from imported data.

Select the **Server Configuration** icon to configure server parameters.

Select the **Preshared Keys** icon to create, delete and edit pre-shared keys.

Select the **Show Public Keys** icon to display the server's public key.

Select the **IPsec Status** icon to display information about the server's capabilities and any current connections..

After a VPN connection is created this menu will include a “Start Connection” button that can start or restart VPN connections. This button is shown in the next view of the VPN Configuration menu.

The "Apply Configuration" button restarts the server to activate any configuration changes that have been made, restarting VPN connections.

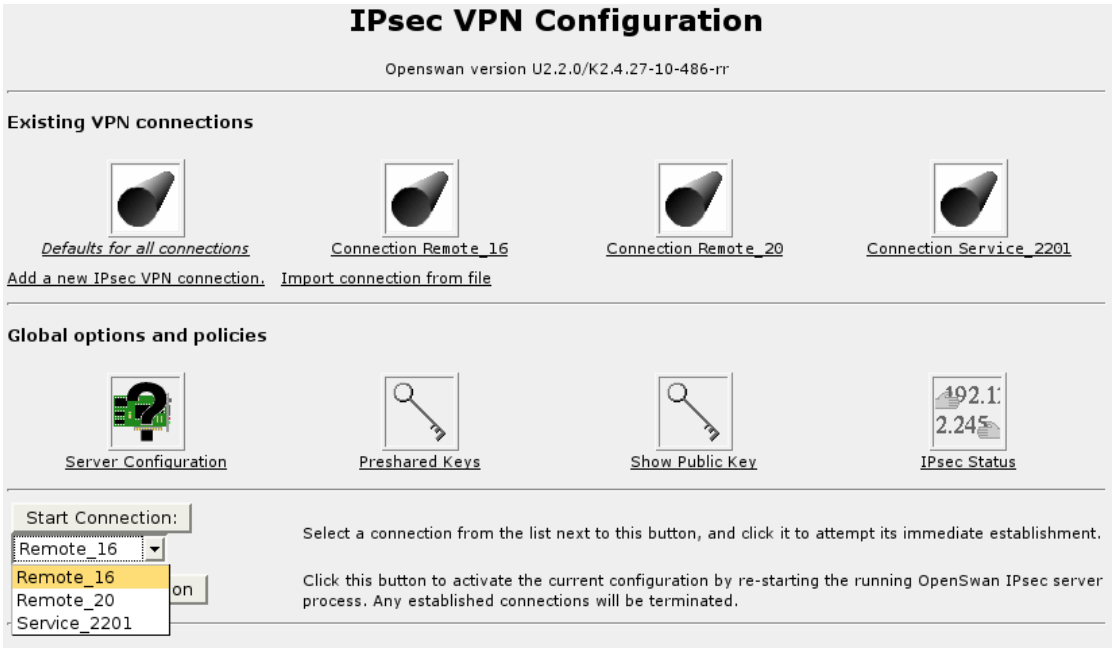

Figure 106: IPsec VPN Configuration After Connections Have Been Created

Server Configuration

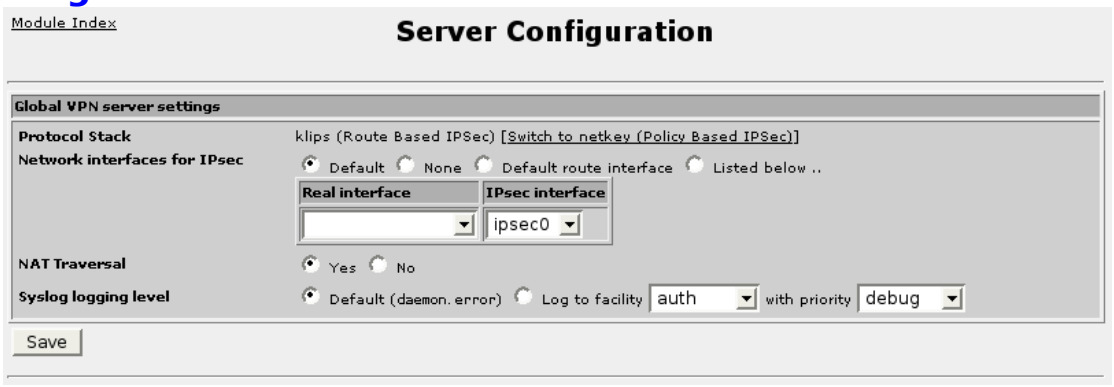

Figure 107: Server Configuration

The **Protocol Stack** field configures whether route based or policy based VPNs are used. Following the link will take the user to menu that requests a reconfirmation and then changes the style of VPN.

The **Network interfaces for IPsec** table configures the association between ipsec interfaces and the real interfaces upon which they become available.

If the **Default** field is selected, Openswan will use its current default (Default route interface at the time of writing) to associate the named ipsec interface with.

If the **Default route interface** field is selected, Openswan will use the real interface owning the default route to associate the named ipsec interface with.

If the **Default** field is selected, Openswan will use its current default (Default route interface at the time of writing) to associate the named ipsec interface with.

If the **Listed below..** field is selected, Openswan will establish the real to ipsec interfaces listed.

Note: *When connections become active, Openswan assigns them to ipsec interfaces. You must plan on these interfaces being the source of incoming traffic in firewall rules.*

The **NAT Traversal** fields enable and disable this feature. Enable Nat Traversal if this router originates the VPN connection and the VPN connection passes through a firewall.

The **Syslog logging level** fields determines the facility and priority of log messages generated by Openswan.

Public Key

Figure 108: Show Public Key

This menu displays the RuggedRouters's public RSA key.

Preshared Keys

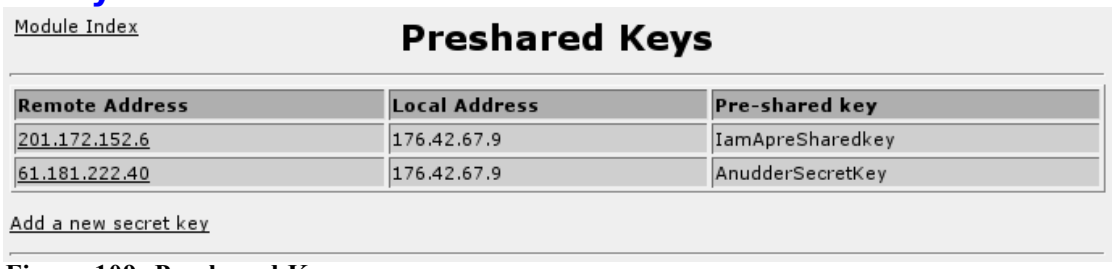

Figure 109: Preshared Keys

This menu creates, deletes and edits pre-shared keys used by VPN connections using secret key encryption.

Select the links under the “Remote Address” column to edit or delete a secret key.

The menu will not allow more than one entry to have a specific pair of IP addresses. The menu will not allow a password shorter than eight characters in length.

List Certificates

Figure 110: List Certificates

This menu lists available certificate files, their corresponding key files and details whether a public key for the certificate is configured.

VPN Connections

The IPsec main menu "Add a new IPsec VPN connection" link leads to the "Create Connection" menu, creating a new connection and its icon. Selecting the connection's icon from the IPsec main menu displays the same menu, allowing editing and deletion.

An IPsec connection is composed of three types of information. There is information about the the local host, the remote host and about the overall connection between them. The configuration data has been designed in such a way that there are identical connection specifications on both ends. Because of this, connection specifications are written in terms of "left" and "right" participants, rather than in terms of local and remote. Which participant is considered left or right is arbitrary; IPsec figures out which one it is being run on based on internal information.

The Create/Edit Connection menu is reflects this organization by being split into three sections. The first section (IPsec VPN Connection Details) describes parameters relating to the connection itself.

The next two sections (Left System's Settings, Right System's Settings) describe IP networking parameters and RSA signatures at each peer. These two sections are identical and are described once.

IPsec VPN Connection Details

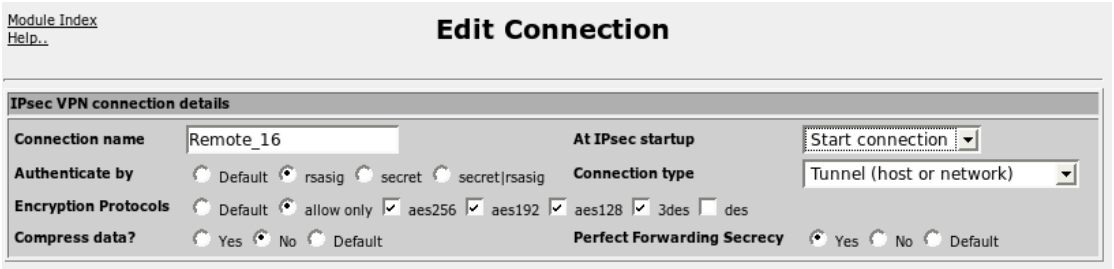

Figure 111: Editing A VPN Connection, Part 1

The **Connection name** field associates a name with the connection. Do not embed whitespace in the name.

The **At IPsec startup** field determines what happens to the connection after Openswan starts and includes the options “Ignore”, “Add connection”, “Start Connection”, “Route” and “Default”. A value of “Ignore” will cause the connection to be ignored. A value of “Add connection” will cause the connection to be established when explicitly started (via command line or the **IPsec VPN Configuration** menu “Start Connection” button). If “Start connection” is chosen then the connection will be authorized when Openswan is started, but not activated until an incoming request arrives. A value of “Route” will cause a route (and only the route) for packets to be established, discarding packets sent there, which may be preferable to having them sent elsewhere based on a more general route (e.g., a default route).

The **Authenticate by** fields select the authentication method. If “Default” is selected the value in the “*Defaults for all connections”* record is used. If “rsasig” or “secret|rsasig” is selected then the **System's public key** of each of the Left System's Settings and Right System's Settings sections must include an RSA signature string or an X.509 certificate must be in use. If “secret” is selected then the **Preshared key** menu must contain a key indexed by the Public IPs of the Left and Right systems.

The **Encryption Protocols** fields select the encryption protocol used. If “Default” is selected the value in the “*Defaults for all connections”* record is used. If “allow only” is selected, the protocols in “aes256”, “aes192”, “aes128” and “3des”, are included in a list. At connection time the two peers will compare their capabilities and select the strongest common protocol (largest aes over smaller aes and aes over 3des).

The **Compress data?** fields will select whether data should be compressed. If “Default” is selected the value in the “*Defaults for all connections”* record is used.

The **Perfect Forward Secrecy** fields will enable PFS, causing keys to be exchanged in a manner which provides attackers that have compromised a key with no advantage in decoding previously intercepted packets or with subsequent packets. Not all clients support PFS.

Left/Right System's Settings

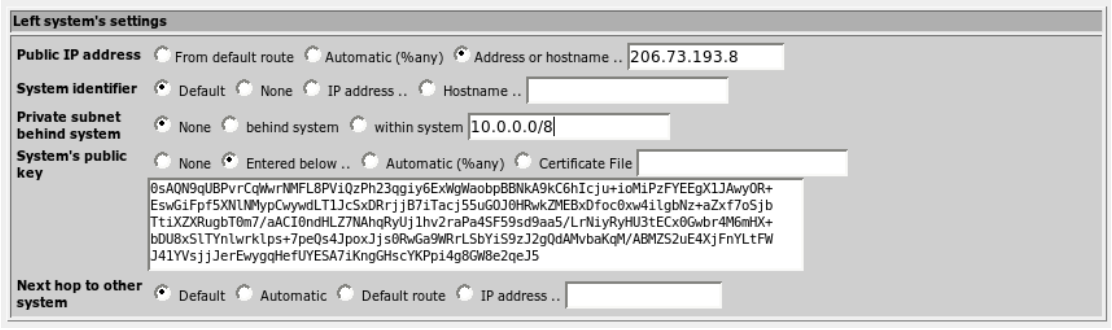

Figure 112: Editing A VPN Connection, Part 2

The **Public IP address** fields determine the IP address of the side of the connection being edited. Check the **Address or hostname.. field** and provide a fixed IP address or hostname. If this side reflects a remote client whose IP address changes, select **Automatic (%any).** Use **From default route** if the host's IP is dynamically assigned.

The **System identifier** fields provide IPsec with a way to determine which section of the connection applies to which host. Left to **Default** the parameter will use the public IP address from above. Set to **None,** the router will use an empty id. You can override these with an IP address or hostname.

The **Private subnet behind system** fields determine if this system has an internal network connected to it that the other host should be granted access to. Enter a network address and prefix length into this field. If you enter a subnet of 0.0.0.0/0 in this field, this connection will serve as a default route for all traffic.

The **System's public key** fields provide an RSA key if RSA keying is to be used. If you want to use secret keying, select **None**. When you first create a connection, this field is filled in for you with the local system's RSA key. If you are filling in this field for the remote system, the key can be obtained from the **Show Public Key** page on that system. Select **Certificate File** and provide a certificate if using X.509 certificates.

The **Next hop to other system** fields determine the address to forward traffic to in order to reach the other system. Unless you have an unusual network setup, this field should be set to **Default route**.

Note: *If you set* ***Next hop to other system*** *to "default", you must configure a default route. You can check for the existence of a default route with the Network Configuration menu, Current Routing & Interface Table icon . A default route will be indicated by a "default" in the Destination column.*

Export Configuration

Selecting the "Export Configuration" button provides a means to capture the connection specification in such a way as to be importable at the remote router.

Showing IPsec Status

```
1  interface lo/lo 127.0.0.1
2  interface eth1/eth1 10.0.0.253
3  interface eth2/eth2 204.50.190.89
4  interface w1ppp/w1ppp 206.186.238.138
5  %myid = (none)
6  debug none

7  algorithm ESP encrypt: id=2, name=ESP_DES, ivlen=8, keysizemin=64, keysizemax=64
8  algorithm ESP encrypt: id=3, name=ESP_3DES, ivlen=8, keysizemin=192, keysizemax=192
9  algorithm ESP encrypt: id=7, name=ESP_BLOWFISH, ivlen=8, keysizemin=40, keysizemax=448
10 algorithm ESP encrypt: id=11, name=ESP_NULL, ivlen=0, keysizemin=0, keysizemax=0
11 algorithm ESP encrypt: id=12, name=ESP_AES, ivlen=8, keysizemin=128, keysizemax=256
12 algorithm ESP encrypt: id=252, name=ESP_SERPENT, ivlen=8, keysizemin=128, keysizemax=256
13 algorithm ESP encrypt: id=253, name=ESP_TWOFISH, ivlen=8, keysizemin=128, keysizemax=256
14 algorithm ESP auth attr: id=1, name=AUTH_ALGORITHM_HMAC_MD5, keysizemin=128, keysizemax=128
15 algorithm ESP auth attr: id=2, name=AUTH_ALGORITHM_HMAC_SHA1, keysizemin=160, keysizemax=160
16 algorithm ESP auth attr: id=5, name=AUTH_ALGORITHM_HMAC_SHA2_256, keysizemin=256, keysizemax=256
17 algorithm ESP auth attr: id=251, name=(null), keysizemin=0, keysizemax=0

18 algorithm IKE encrypt: id=7, name=OAKLEY_AES_CBC, blocksize=16, keydeflen=128
19 algorithm IKE encrypt: id=5, name=OAKLEY_3DES_CBC, blocksize=8, keydeflen=192
20 algorithm IKE hash: id=2, name=OAKLEY_SHA, hashsize=20
21 algorithm IKE hash: id=1, name=OAKLEY_MD5, hashsize=16
22 algorithm IKE dh group: id=2, name=OAKLEY_GROUP_MODP1024, bits=1024
23 algorithm IKE dh group: id=5, name=OAKLEY_GROUP_MODP1536, bits=1536
24 algorithm IKE dh group: id=14, name=OAKLEY_GROUP_MODP2048, bits=2048
25 algorithm IKE dh group: id=15, name=OAKLEY_GROUP_MODP3072, bits=3072
26 algorithm IKE dh group: id=16, name=OAKLEY_GROUP_MODP4096, bits=4096
27 algorithm IKE dh group: id=17, name=OAKLEY_GROUP_MODP6144, bits=6144
28 algorithm IKE dh group: id=18, name=OAKLEY_GROUP_MODP8192, bits=8192

29 stats db_ops.c: {curr_cnt, total_cnt, maxsz} :context={0,6144,36} trans={0,6144,336} attrs={0,6144,224}

30 "openswantest": 10.0.0.0/8===204.50.190.89...204.50.190.91===192.168.1.0/24; erouted; eroute owner: #2997
31 "openswantest":   ike_life: 3600s; ipsec_life: 28800s; rekey_margin: 540s; rekey_fuzz: 100%; keyingtries: 0
32 "openswantest":   policy: PSK+ENCRYPT+TUNNEL+PFS+UP; prio: 24,8; interface: eth2;
33 "openswantest":   newest ISAKMP SA: #3093; newest IPsec SA: #2997;
34 "openswantest":   IKE algorithms wanted: 5_000-1-5, 5_000-1-2, 5_000-2-5, 5_000-2-2, flags=-strict
35 "openswantest":   IKE algorithms found:  5_192-1_128-5, 5_192-1_128-2, 5_192-2_160-5, 5_192-2_160-2,
36 "openswantest":   IKE algorithm newest: 3DES_CBC_192-MD5-MODP1536
37 "openswantest":   ESP algorithms wanted: 3_000-1, 3_000-2, flags=-strict
38 "openswantest":   ESP algorithms loaded: 3_000-1, 3_000-2, flags=-strict
39 "openswantest":   ESP algorithm newest: AES_256-HMAC_SHA1; pfsgroup=<Phase1>

40 #3126: "openswantest" STATE_QUICK_I1 (sent QI1, expecting QR1); EVENT_RETRANSMIT in 9s
41 #3093: "openswantest" STATE_MAIN_R3 (sent MR3, ISAKMP SA established); EVENT_SA_REPLACE in 1050s; newest ISAKMP
42 #2997: "openswantest" STATE_QUICK_R2 (IPsec SA established); EVENT_SA_REPLACE in 19773s; newest IPSEC; eroute owner
43 #2997: "openswantest" esp.df9839e9@204.50.190.91 esp.8e2d7255@204.50.190.89 tun.0@204.50.190.91 tun.0@204.50.190.
```

Figure 113: IPsec Status

The “IPsec Status” button produces a window of text similar to that of the above figure (except that line numbers have been inserted for purposes of illustration).

The first group (lines 1-5) describes configured interfaces.

The second group (lines 7-17) describes ESP capabilities. In this group we can see encryption capabilities (lines 7-13) and authentication capabilities (lines 14-17). At least one set of values must match between the left- and right-hand side VPN devices. This is also frequently referred to as the Phase 2 parameters, because the data encryption process is the second and final thing to occur in establishing a VPN.

The third group (lines 18-28) describes IKE capabilities and defines the various encrypted key exchange algorithms and their parameters. At least one set of values must match between the left- and right-hand side VPN devices. This is also frequently referred to as the Phase 1 parameters, because the key exchange process is the first thing to occur in establishing a VPN.

The fourth group (lines 30-39) describe connection describe VPN connections (here "openswantest"). The first line is particularly useful since it indicates the connection addresses, subnets and that the connection is active ("erouted"). If there are no entries, then the VPN hasn't been established at all. If there are entries, but no STATE_QUICK_R2 (IPsec SA established) lines then the IPSec parameters are configured, but the tunnel hasn't been established. This can be normal, tunnels become active once the Phase 1 and Phase 2 security associations are created, and this usually only occurs after traffic is flowing. The associations then get torn down after a timeout period.

IPSec X.509 Roaming Client Example

This example details how to set up IPSec connections using X.509 certificates on the router. The router will provide an IPSec gateway to a number of remote clients that connect via an Internet connection. Each of the clients will fetch an IP address locally from a DHCP server, and it is assumed (but not required) that network address translation will be applied at the client end. Each of the clients should "appear" on the local network on a specific IP address. In this example the clients are laptop PCs.

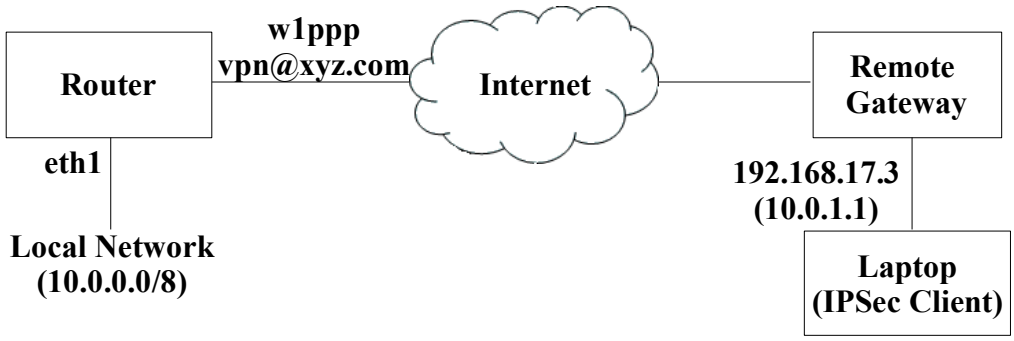

Figure 114: End To End Backup Example

Select A Certificate Authority

Begin by constructing the required certificates. You may construct the certificates using a RuggedRouter or a third party tool. The device that is used to build the certificates is known as the certificate authority. There are advantages and disadvantages to using the router itself as the authority. It is convenient to use if it is the only router in the network and many clients will be connecting to it. On the other hand, if the router holds the certificate authority and is compromised, all certificates must be constructed again.

Ensure that the the Certificate Authority generates certificates with a reasonable life and generates keys of at least 1024 bits in length.

Generate X.509 Certificates

Use the authority to produce a certificate authority public certification (cacert) and a certificate for each of the clients and a certificate for the router. The certificate authority will require some information that is shared by all certificates (e.g. a Country Name (C), a State Or Province Name (S), an Organization name (O)) and some per-client information (e.g. a Common Name (CN) and an Email address (E)). Together this information forms the Distinguished Name (DN) and is used by the router and client to validate each other.

VPN Networking Parameters

The first step is to identify the key parameters required. The router public gateway (here vpn@xyz.com) and its gateway interface (w1ppp) must be known. The local network subnet (10.0.0.0/8) and each clients' internal network address (here 10.0.1.1) must be known. All client addresses should be assigned from a subnet of the local network (e.g. 10.0.1.0/24). A number of encryption parameters should be decided upon depending upon the client capabilities. Avoid selecting 3DES if possible due to its high overhead.

Client Configuration

Depending upon the client, you may be required to produce the certificate in a P12 format, and may be required to include an "export" password as well. This password will be required to be known be the personnel that configure the client in order to import the certificate.

Install the client IPSec software and import the cacert and the clients own certificate and key. Configure the client with the router public gateway, the clients internal network address and the desired encryption parameters. At this point the client should be able to use its Internet connection to ping the public gateway.

Router IPSec Configuration

Transfer the cacert and the router's certificate to the router. If your authority prepares a Certificate Revocation List (CRL), you will want to transfer that as well.

The cacert file should be renamed cacert.pem and installed in /etc/ipsec.d/cacerts/.

The CRL file should be renamed to crl.pem and installed in /etc/ipsec.d/crls/.

The router's certificate must be installed in /etc/ipsec.d/certs/. It's public key file (e.g. router.key) must be installed in /etc/ipsec.d/private/ and a line ': RSA router.key "Password"' (where Password is the pass phrase that was used to generate the certificate) must be added to the end of the /etc/ipsec.secrets file.

Note: *The Maintenance Menu, Upload/Download Files sub-menu provides a method to transfer the files directly to the indicated directories.*

Enable IPSec from the **Bootup and Shutdown** menu. Visit the **IPSec VPN** menu and generate a public key.

Visit the **Server Configuration** menu and associate the ipsec0 interface with the desired interface the connection will arrive on (here w1ppp).

Create a connection for the clients. Set the parameters as follows:

Parameters	Value	Comments
At IPsec Startup	Add connection	We wish to add the connection when the client starts it.
Authenticate by	rsasig	X.509 certificates provide RSA
Connection Type	Tunnel	
Encryption Protocols	As desired	
Compress Data	As desired	
Perfect Forwarding Secrecy	As desired	Recommend “yes”
NAT Traversal	No	Required when the router acts as a client and is behind a NAT firewall.
Left System Settings		**Router's side**
Public IP Address	Address or hostname .. (IP of public gateway)	
System Identifier	Default	
Private subnet behind system	10.0.0.0/8	
System's public key	Certificate File (router.pem)	
Next hop to other system	Default	
Right System Settings		**Laptop1 side**
Public IP Address	Automatic	
System Identifier	Default	
Private subnet behind system	10.0.1.0/24	Assign IP based on client from within this subnet
System's public key	Entered below (%cert)	Derive identity from incoming certificate
Next hop to other system	Default	

Apply the configuration to restart the server and create an ipsec0 interface.

Firewall IPSec Configuration

Create firewall Zones “vpn” and net. Ensure that the WAN interface (here w1ppp) and ipsec0 interface are present in the Shorewall **Network Interfaces**. The WAN interfaces should be in zone “net” while ipsec0 should be in zone “vpn”.

Add the following firewall rules:

Action	Source-Zone	Destination-Zone	Protocol	Dest-Port
ACCEPT	all	fw	ah	
ACCEPT	all	fw	esp	
ACCEPT	all	fw	udp	500
ACCEPT	vpn	loc		

Restart the firewall to install the rules.

Ethernet Port Configuration

Because the remote client will be assigned a local IP address but is reachable only through the IPSec connection, proxy ARP must be employed. Activate proxy ARP on the Ethernet interface that hosts the local network (here eth1) via the **Networking** Menu, **Ethernet** sub-menu **boot time entry Proxy ARP setting**. When a host on eth1 arps for the remote client address, the router will answer on behalf of the client.

This page intentionally blank

Chapter 13 – Configuring Dynamic Routing

Introduction

This chapter familiarizes the user with:

- Enabling The Dynamic Routing Suite
- Enabling And Starting OSPF and RIP
- Configuring OSPF and RIP
- Obtaining OSPF and RIP Status
- OSPF and VRRP

Quagga, RIP and OSPF

Dynamic routing is provided by the Quagga suite of routing protocol daemons. Quagga provides three daemons for managing routing, the core, ripd and ospfd.

The core daemon handles interfacing with the kernel to maintain the router's routing table and to check link statuses. It tells RIP and OSPF what state links are in, what routes are in the routing table, and some information about the interfaces.

The ripd and ospfd daemon handles communications with other routers using the RIPv2 and OSPFv2 protocol, decides which routers preferred to forward to.

In complex legacy networks, both RIP and OSPF may be active on the same router at the same time. Usually, one on them is employed.

RIP Fundamentals

The Routing Information Protocol determines the best path for routing IP traffic over a TCP/IP network based on the number of hops between any two routers. It uses the shortest route available to a given network as the route to use for sending packets to that network.

The RuggedRouter RIP daemon (ripd) is an RFC1058 compliant implementation of RIP support RIP version 1 and 2. RIP version 1 is limited to obsolete class based networks, while RIP version 2 supports subnet masks as well as simple authentication for controlling which routers to accept route exchanges with.

RIP uses network and neighbor entries to control which routers it will exchange routes with. A network is either a subnet or a physical interface (it must to be a broadcast capable interface). Any router that is part of that subnet or connected to that interface may exchange routes. A neighbor is a specific router to exchange routes with specified by its IP address. For point to point links (T1/E1 links for example) one must use neighbor entries to add other routers to exchange routes with. The maximum number of hops between two points on a RIP network is 15, placing a limit on network size.

Link failures will eventually be noticed although it is not unusual for RIP to take many minutes for a dead route to disappear from the whole network. Large RIP networks could take over an hour to converge when link or route changes occur. For fast convergence and recovery, OSPF is a much better choice. RIP is a fairly old routing protocol and has mostly been superseded by OSPF.

OSPF Fundamentals

The Open Path Shortest First (OSPF) protocol routing determines the best path for routing IP traffic over a TCP/IP network based on link cost and quality. Unlike static routing, OSPF takes link failures and other network topology changes into account. Unlike the RIP routing protocol, OSPF provides less router to router update traffic.

RuggedRouter routing protocols are supplied by the Quaaga routing package.

The RuggedRouter OSPF daemon (ospfd) is an RFC 2178 compliant implementation of OSPFv2. The daemon also adheres to the RFC2370 (Opaque LSA) and RFC3509 (ABR-Types) extensions.

OSPF network design usually involves partitioning a network into a number of self contained areas. The areas are chosen to minimize intra-area router traffic, making more manageable and reducing the number of advertised routes. Area numbers are assigned to each area. All routers in the area are known as Area routers. If traffic must flow between two areas a router with links in each area is selected to be an Area Border router, and serves as a gateway.

Link State Advertisements

When an OSPF configured router starts operating it issues a *hello* packet. Routers having the same OSPF Area, hello-interval and dead-interval timers will communicate with each others and are said to be *neighbors*

After discovering its neighbors, a router will exchange *Link State Advertisements* in order to determine the network topology.

Every 30 minutes (by default) the entire topology of the network must be sent to all routers in an area. If the link speeds are too low, the links too busy or there are too many routes, then some routes may fail to get re-announced and will be aged out.

Splitting the network into smaller areas to reduce the number of routes within an area or reducing the number of routes to be advertised may help to avoid this problem.

In shared access networks (i.e. routers connected by switches or hubs) a designated router and a backup designated are elected to receive route changes from subnets in the area. Once a designated router is picked, all routing state changes are sent to the designated router, which then sends the resulting changes to all the routers.

The election is decided based on the priority assigned to the interface of each router. The highest priority wins. If the priority is tied, the highest router-id wins.

Key OSPF And RIP Parameters

Network Areas

Network areas determine the regions within which routes are distributed to other routers. The subnets at a particular router can be added to its OSPF Area. The router will advertise these subnets to all routers in its area.

Note: *OSPF areas must be designed such that no single link failure will cause the network to be split into two disjoint networks.*

A router can be part of multiple areas and function as a gateway between areas. When multiple areas are used on a network, area 0 is the backbone area. All areas must have a router connecting them to area 0.

Router-ID

Defines the ID of the router. By default this is the highest IP assigned to the router. It is often a good idea to configure this value manually to avoid the router-id changing if interfaces are added or deleted from the router. During elections for designated router, the router-id is one of the values used to pick the winner. Keeping the router-id fixed will avoid any unexpected changes in the election of the master router.

Hello Interval and Dead Interval

The hello interval is the time between transmission of OSPF Hello packets. The dead interval is the time to wait without seeing an OSPF Hello packet before declaring a neighboring router dead and discarding its routes. It is recommended that the dead interval be at least four times the hello interval for reliable operation.

Lower values of these settings will help to speed up the change in network routes when the topology of the network changes. It will also increase the load on the router and the links, due to higher traffic caused by the increase in messages. Lower values will also put limits on the number of routes that can be distributed within an area, as will running over slower links.

Note: *OSPF will not work properly if the Hello Interval and Dead Interval are not identical on every router in an area.*

Active/Passive Interface Default

OSPF regards router interfaces as either passive or active, sending OSPF messages on active interfaces and ignoring passive interfaces. By default, newly created interfaces are viewed as passive from OSPF until they are configured active. This is more efficient and secure for the router. The default type for new interfaces is controlled by the passive interface default option in the OSPF Global Parameters.

Note: *The default setting of Passive Interface Default means that you must explicitly configure interfaces active before OSPF will attempt to use them.*

Redistributing Routes

Routes for subnets which are directly connected to the router but are not part of the OSPF area or RIP network can be advertised if “redistribute connected” is enabled in the OSPF or RIP Global Parameters. Static routes and other routes handled by the kernel can also be redistributed if redistribute kernel is enabled.

Link Detect

When link detect is enabled for an OSPF/RIP active interface, OSPF or RIP will be notified when the interface goes down and will stop advertising subnets associated with that interface. OSPF and RIP will resume advertising the subnet when the link is restored. This allows OSPF and RIP to detect link failures more rapidly (as the router does not have to wait a dead interval to time out). Link Detect will also cause " redistributed" routes to start and stop being advertised based upon the status of their interface links.

Configuring OSPF Link Costs

Link cost is used when multiple links can reach a given destination, to determine which route to use. OSPF will (by default) assign the same cost to all links unless provided with extra information about the links. Each interface is assumed to be 10Mbit unless told otherwise in the Core Interface configuration.

The reference bandwidth for link cost calculations is 100Mbit by default in the OSPF Global Parameters. The reference bandwidth divided by the link bandwidth gives the default cost for a link, which by default is 10. If a specific bandwidth is assigned to each link, the costs will take this into account.

It is also possible to manually assign a cost to using a link in the OSPF Interface Configuration for each interface for cases where the speed of the link is not desired as the method for choosing the best link.

OSPF Authentication

OSPF authentication is used when it is desirable to prevent unauthorized routers from joining the OSPF network. By enabling authentication and configuring a shared key on all the routers, only routers which have the same authentication key will be able to send and receive advertisements within the OSPF network. Authentication adds a small overhead due to the encryption of messages, so is not to be preferred on completely private networks with controlled access.

RIP Authentication

RIP authentication is used when it is desirable to prevent unauthorized routers from joining the network. RIP authentication is supported by per-interface configuration or the use of key-chains. Separate key chains spanning different groups of interfaces and having separate lifespans are possible. By enabling authentication and configuring a shared key on all the routers, only routers which have the same authentication key will be able to send and receive advertisements within the RIP network.

OSPF And Antispoofing

Antispoofing is the process of discarding packets arriving on an interface because they match the subnet of another configure interface. This is not a normal occurrence in conventional routing. This situation can arise in OSPF, when routers are multiply connected. If for example two routers are connected by lower speed wan and higher speed Ethernet links, packets on subnets native to the wan will still be forwarded via Ethernet because of cost. If antispoofing is enabled, the packet will be discarded at the peer OSPF router.

Note: *Ensure that Antispoofing is disabled if you are constructing the above described type of OSPF network. Antispoofing can be disabled in the* ***Network Configuration*** *menu,* ***Core Settings*** *sub-menu.*

Administrative Distances

The router may work with different routing protocols at the same time, as well as employing local interface and statically assigned routes. An administrative distance, (from 0 to 255) is a rating of the trustworthiness of a routing information source. For a given route, the protocol having the lowest administrative distance will be chosen. By default the distances for a connected interface is 0 and for a static route is 1. By default, OSPF will set an administrative distance of 110 and RIP will set a distance of 120.

OSPF And VRRP Example Network

This network consists of three routers connected in a ring with T1/E1 links. Router 1 and 2 and the switched network represent a remote site in which the routers supply a redundant gateway to the hosts via VRRP and the T1/E1 links supply a redundant network connection to the rest of the network.

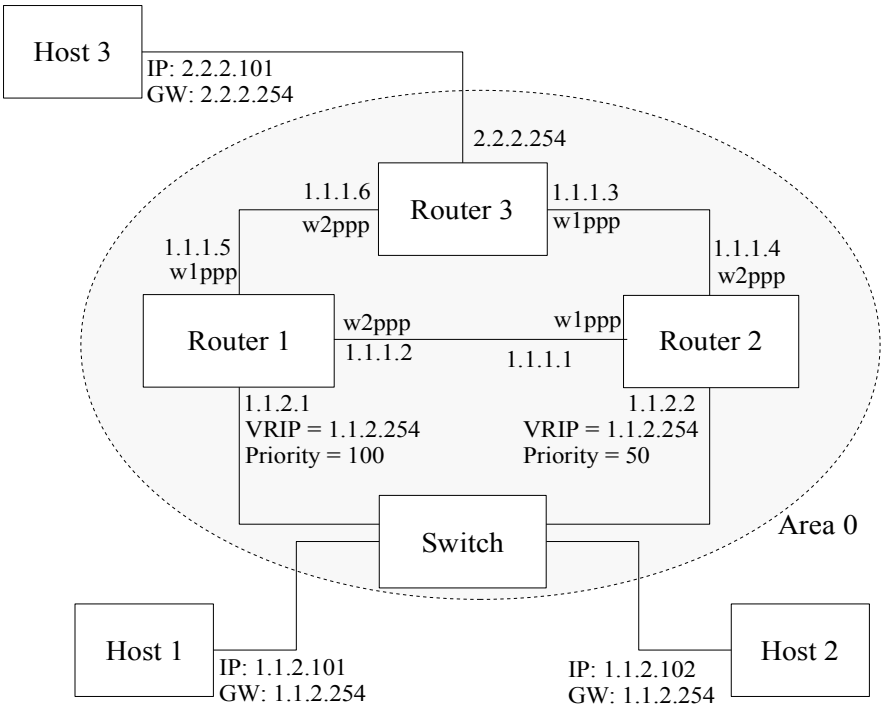

Figure 115: OSPF And VRRP Example

Area And Subnets

As the OSPF design is simple, an area of 0 is used. The three point-to-point T1/E1 links are placed in the area by adding 1.1.1.0/24 to it. Router 1 and 2 will include their Ethernet links by adding subnet 1.1.2.0/24 to their area descriptions. Router 3 must also include 2.2.2.0/24 in its area description so that its existence is advertised.

The point-to-point T1/E1 interfaces and Ethernet interfaces on Router 1 and 2 must be made active. The Ethernet interface on Router 3 can be left passive since it does not participate in OSPF advertisements.

Router 1 and 2 must enable link-detect, to stop advertising 1.1.1.0/24 in the event of a link failure.

VRRP Operation

Router 1 and 2 have VRRP setup on their Ethernet connection so that they can both function as the gateway for the clients on their network segment. Normally Router 1 is the VRRP master, and only in case of a link failure to the switch or the router failing, will Router 2 take over the virtual IP. The virtual IP used as the gateway is 1.1.2.254. Each router also has its own IP on the network so that each can be reached individually.

If Router 1 or its Ethernet link fail, VRRP will detect the link being down and remove the direct route to the 1.1.2.0/24. VRRP on Router 2 will stop seeing messages from Router 1, elect itself master and will take over the gateway for the network.

OSPF on router 1 will notice the link being down (and the route to 1.1.2.0/24 disappearing) and will use information from router 2 install a route to 1.1.2.0/24 via Router 2.

Router 3 will notice than Router 2 is now a more direct path to 1.1.2.0/24 network and start sending to Router 2 instead of Router 1.

After the failure all routers still know how to reach the entire network, and the clients on 1.1.2.0/24 can still send on the network using the same gateway address. The clients will see only a MAC address change of the gateway and experience a few seconds of network outage When the link returns, VRRP will switch back to the master, and the routes will return to their normal state.

Note that if the Router 1 wan link fails, Router will see routes to Router3 via the Router 1 – Router 2 wan and Ethernet links. If the faster Router 1 – Router 2 Ethernet path fails, Router 1 will fall back to the Router 1 – Router 2 wan link.

Note that it would not be useful to leave the Ethernet 1.1.2.0/24 subnets out of the area and turn on redistribute connected as OSPF would not use the subnets for routing.

Dynamic Routing

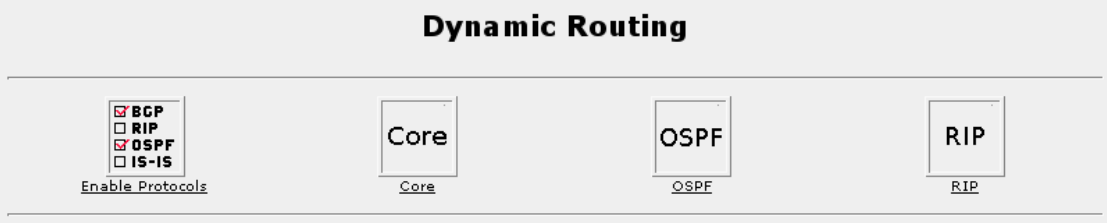

Figure 116: Dynamic Routing Menu

Before dynamic routing protocols can be used, quagga must be enabled in the Bootup and Shutdown menu.

After quagga is enabled, RIP or OSPF itself must be enabled in the Enable Protocols menu of Dynamic Routing.

The Core menu configures link related items such as link-detect and link cost.

The RIP and OSPF menu configure these protocols for each interface.

Enable Protocols

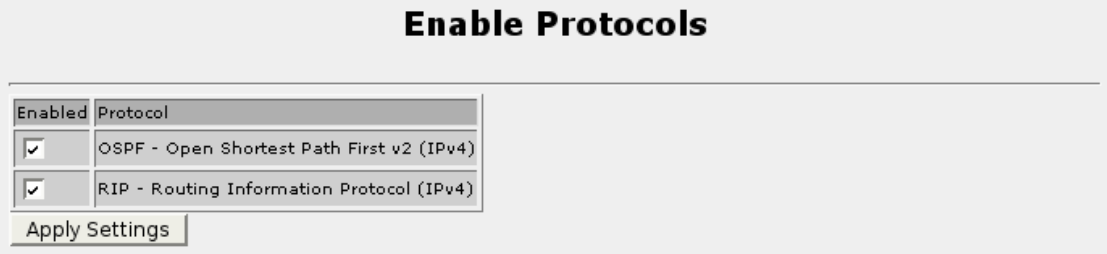

Figure 117: Enable Protocols Menu

This menu enables RIP and OSPF for dynamic routing.

Core

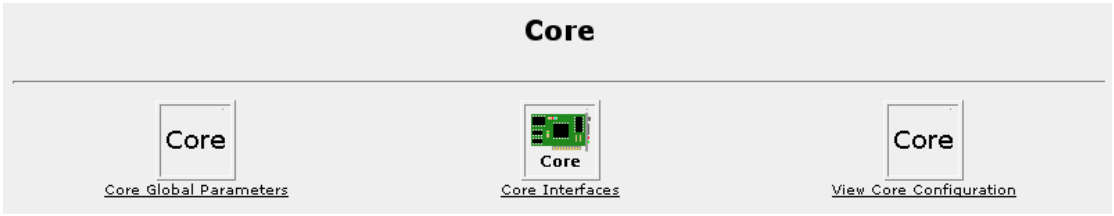

Figure 118: Core Menu

The Core routing daemon handles communications between the kernel of the router and the other dynamic routing protocols. Core handles link detection and monitoring static routes and routes for directly connected interfaces on the router. It also manages adding routes to the kernel routing table based on the routes discovered by other dynamic routing protocols. Core is always enabled whenever dynamic routing is enabled as it is required by all other dynamic routing protocols.

Core Global Parameters

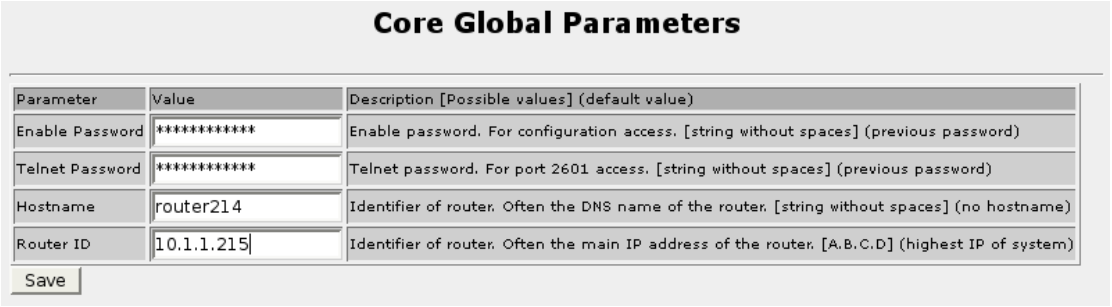

Figure 119: Core Global Parameters

The **Enable Password** field sets the password to be used for the enable command of core. This is used by the telnet interface of core to control access to the configuration.

The **Telnet Password** field sets the password to be used for telnet access to core. This is used as the login password of core when locally telnetting to port 2601 of the router.

The **Hostname** field sets the hostname for the core daemon. This value is only used as a a reference for convenience. The telnet interface prompt will contain this hostname. The router's system wide hostname is used if this field is left blank.

The **Router Id** field sets the router-id to use for the core daemon. This value is used as a unique identifier for the dynamic routing protocol to identify which router sent which route advertisement. By default it uses the highest IP assigned to an interface on the router. It is recommended that this value be set to a unique fixed IP on each router.

Core Interface Parameters

Figure 120: Core Interface Parameters

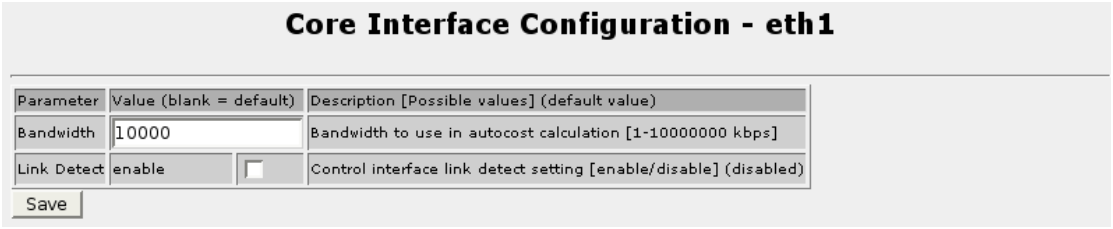

Parameters specific to one interface are configured here.

Each interface on the router is listed. Clicking on settings displays a menu of configuration options for that interface. Clicking on status displays the current status of the interface, including link state, IP address and traffic counts.

Clicking “Remove inactive interfaces” purges the list of any interfaces which are no longer configured on the router.

The **Bandwidth** field sets the bandwidth value to assume for the interface when automatically calculating a cost for using the link on this interface. By default all interfaces are treated as 10Mbit (10000 Kbps). OSPF by default uses an automatic cost of 10 for all links by calculating is as reference bandwidth (100Mbit) divided by the link bandwidth (10Mbit). If a manual cost is assigned to the interface in OSPF, this value is ignored. RIP does not use this parameter.

The **Link Detect** field controls core's link detect feature on the interface. When link detect is enabled, routes through the interface will only be advertised to other routers when the link is up. This option is usually desirable.

View Core Configuration

This menu shows the current configuration file for the Core interfaces.

OSPF

Figure 121: OSPF Menu

This menu contains the configuration and status of OSPF on the router.

The **OSPF Global Parameters**, **OSPF Interfaces** and **Network Areas menus** configure OSPF. The Status and View OSPF Configuration menu display the actual status and configuration file contents of OSPF.

OSPF Global Parameters

OSPF Global Parameters

Parameter	Value	Description [Possible values] (default value)
Enable Password	**********	Enable password. For configuration access. [string without spaces] (previous password)
Telnet Password	**********	Telnet password. For port 2604 access. [string without spaces] (previous password)
ABR-Type	standard	Set OSPF ABR type [standard/cisco/ibm/shortcut] (standard)
Auto Cost Reference Bandwidth	100	Calculate OSPF interface cost according to bandwidth [1-4294967 Mbps] (100)
Default-Information Originate	enable ☐	Advertise default route (disabled)
Default Metric	20	Control distribution of default information [1-16777214] (20)
Distance		Define an administrative distance [unset,1-255] (unset=not used)
Distance OSPF External		Define an administrative distance (external) [unset,1-255] (unset=use Distance)
Distance OSPF Inter-area		Define an administrative distance (inter-area) [unset,1-255] (unset=use Distance)
Distance OSPF Intra-area		Define an administrative distance (intra-area) [unset,1-255] (unset=use Distance)
Hostname		Identifier of router. Often the DNS name of the router. [string without spaces] (no hostname)
Opaque LSA	enable ☐	Enable Opaque LSA capability (disabled)
Passive Default	enable ☐	Set new interfaces passive by default (enabled)
Refresh Timer	10	Set refresh timer [10-1800 Seconds] (10)
RFC 1583 Compatibility	enable ☐	Enable compatibility with obsolete RFC1583 OSPF (current is RFC2178) (disabled)
Redistribute Connected	enable ☐ metric-type 2 metric	Redistribute routes for directly connected interfaces to OSPF area routers. [enable/disable,1/2,0-16777214] (disabled,2,unset)
Redistribute Kernel	enable ☐ metric-type 2 metric	Redistribute kernel routes to OSPF area routers. [enable/disable,1/2,0-16777214] (disabled,2,unset)
Redistribute RIP	enable ☐ metric-type 2 metric	Redistribute rip routes to OSPF area routers. [enable/disable,1/2,0-16777214] (disabled,2,unset)
Router ID		Identifier of router. Often the main IP address of the router. [A.B.C.D] (highest IP of system)

Save

Figure 122: OSPF Global Parameters

The **Enable Password** field sets the password to be used for the enable command of ospfd. This is used by the telnet interface of ospfd to control access to the configuration.

The **Telnet Password** field sets the password to be used for telnet access to ospfd. This is used as the login password of ospfd when locally telnetting to port 2604 of the router.

The **ABR-Type** field select which method to use on area border routers to manage inter area routes. Standard follows RFC2178, Cisco and IBM follow RFC3509. Shortcut is covered by the draft-ietf-ospf-shortcut-abr-00.txt document. Standard requires all ABRs to have a backbone connection. The other three methods allow for ABRs that do not have a backbone connection.

The **Auto Cost Reference Bandwidth** field sets the reference bandwidth used to calculate auto costs for OSPF interfaces. The auto cost is the reference bandwidth divided by the interface bandwidth. By default this is 100Mbit/10Mbit = auto cost of 10. The interface cost is set in the Core Interface configuration for each interface. The cost for each interface can also be set in the OSPF Interface configuration to override the auto cost calculation.

The **Default Metric** field sets the default metric to be used for OSPF routes which don't have another metric specified.

The **Default-Information Originate** field, when enabled, causes the router to advertise its default route to the OSPF network.

The **Distance** field sets the administrative distance to use for all routes unless overridden by other distance settings.

The **Distance External** field sets the administrative distance to use for all external routes (backbone routes). The **Distance Inter-area** field sets the administrative distance to use for all routes between areas. The **Distance Intra-area** field sets the administrative distance to use for all routes within an area.

The **Hostname** field sets the hostname for the ospf daemon. This value is only used as a a reference for convenience. The telnet interface prompt will contain this hostname. The router's system wide hostname is used if this field is left blank.

The **Opaque LSA** field controls the opaque LSA option. This feature is covered in RFC2370. This feature is sometimes used to distribute application specific information through a network using OSPF LSAs.

The **Passive Default** option controls the default active/passive state of new interfaces. When enabled all new interfaces will be passive by default. The passive state of individual interfaces is controlled from the OSPF Interfaces configuration.

The **Refresh Timer** field controls how frequently OSPF LSA refreshes occur.

The **RFC 1583 Compatibility** field controls support for RFC1583 compatibility. If this option is enabled OSPF will be compatible with the obsolete RFC1583 version of OSPF. By default it is compatible with RFC2178 version of OSPF only.

The **Redistribute Connected** fields control distribution of connected routes. When enabled, OSPF will advertise routes to directly connected interfaces to other OSPF routers in the area. Normally only routes that fall within the scope of the network areas will be advertised.

The **Redistribute Kernel** fields control distribution of kernel routes. When enabled, OSPF will advertise routes from the kernel routing table, which includes static routes entered by the administrator, to other OSPF routers in the area. Normally only routes that fall within the scope of the network areas will be advertised.

The **Redistribute RIP** fields control distribution of routes learned by RIP. When enabled, OSPF will advertise routes learned by RIP.

The **Router Id** field sets the router-id to use for the ospf daemon. This value is used as a unique identifier for the dynamic routing protocol to identify which router sent which route advertisement. By default it uses the highest IP assigned to an interface on the router. It is recommended that this value be set to a unique fixed IP on each router.

OSPF Interfaces

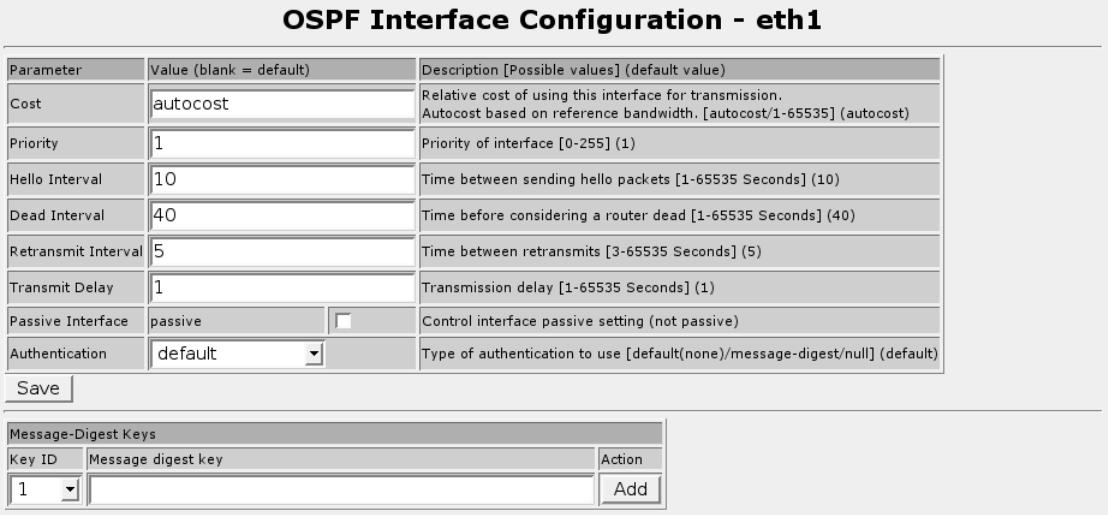

Figure 123: OSPF Interfaces

Parameters specific to one interface are configured here.

Each interface on the router is listed. Clicking on settings displays a menu of configuration options for that interface. Clicking on status displays the current status of the interface, including link state and current OSPF status on the interface. If an interface is not part of an area it will show up as OSPF not enabled on interface.

Clicking “Remove inactive interfaces” purges the list of any interfaces which are no longer configured on the router.

The **Cost** field controls the administrative cost of routing over this interface. By default the cost is auto calculated as the ospf reference bandwidth divided by the core interface bandwidth. By default this is 100Mbit/10Mbit = cost 10.

The **Priority** field controls the priority associated with this interface. By default the priority of interfaces is 1. The router with the highest priority wins elections for designated router for an area.

The **Hello Interval** field controls how often hello packets are sent to other routers in the area. This value must match on all router interfaces in an area.

The **Dead Interval** field controls how long to wait for hello packets before declaring another router dead. This should normally be set to 4 times the hello interval.

The **Retransmit Interval** field controls the delay between retransmissions.

The **Transmit Delay** field controls the estimated number of seconds to transmit a link state update packet. This should take into account transmission and propagation delays of the interface.

The **Passive Interface** option controls if an interface is active or passive. Passive interfaces do not send LSAs to other routers and are not part of an OSPF area.

The **Authentication** field controls the type of authentication to use when communicating with other routers. It can be none, null (just check for message corruption) or message digest, which cryptographically signs each message with a shared key.

The Message Digest Keys fields allows for addition and deletion of keys to use for areas connected to this interface when authentication is set to message-digest.

OSPF Network Areas

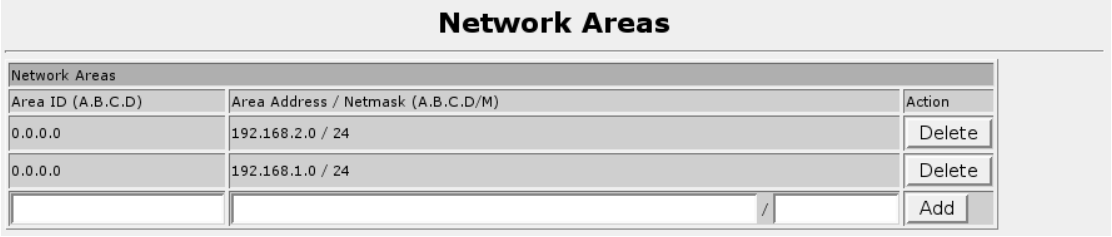

Figure 124: Network Areas

OSPF uses areas to control which routes are distributed between routers. To add a network to an area, enter the area id and the network address and netmask and click Add. To delete an entry click the Delete button beside the entry. All networks routes that are part of the same area will be distributed to other routers in the same area.

OSPF Status

This status menu shows various pieces of information about the current OSPF status. The status of each interface is shown, the current database, the current OSPF neighbors and the current OSPF routing table.

View OSPF Configuration

This menu shows the current configuration file of OSPF.

RIP

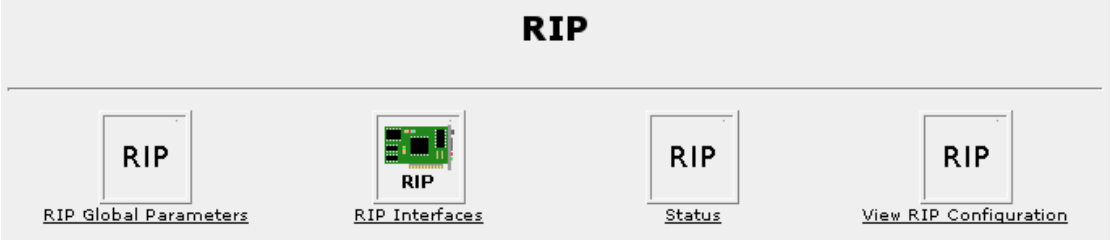

Figure 125: RIP Menu

This menu contains the configuration and status of RIP on the router.

The **RIP Global Parameters and RIP Interfaces** configure RIP. The **Status** and **View RIP Configuration** menu display the actual status and configuration file contents of RIP.

RIP Global Parameters

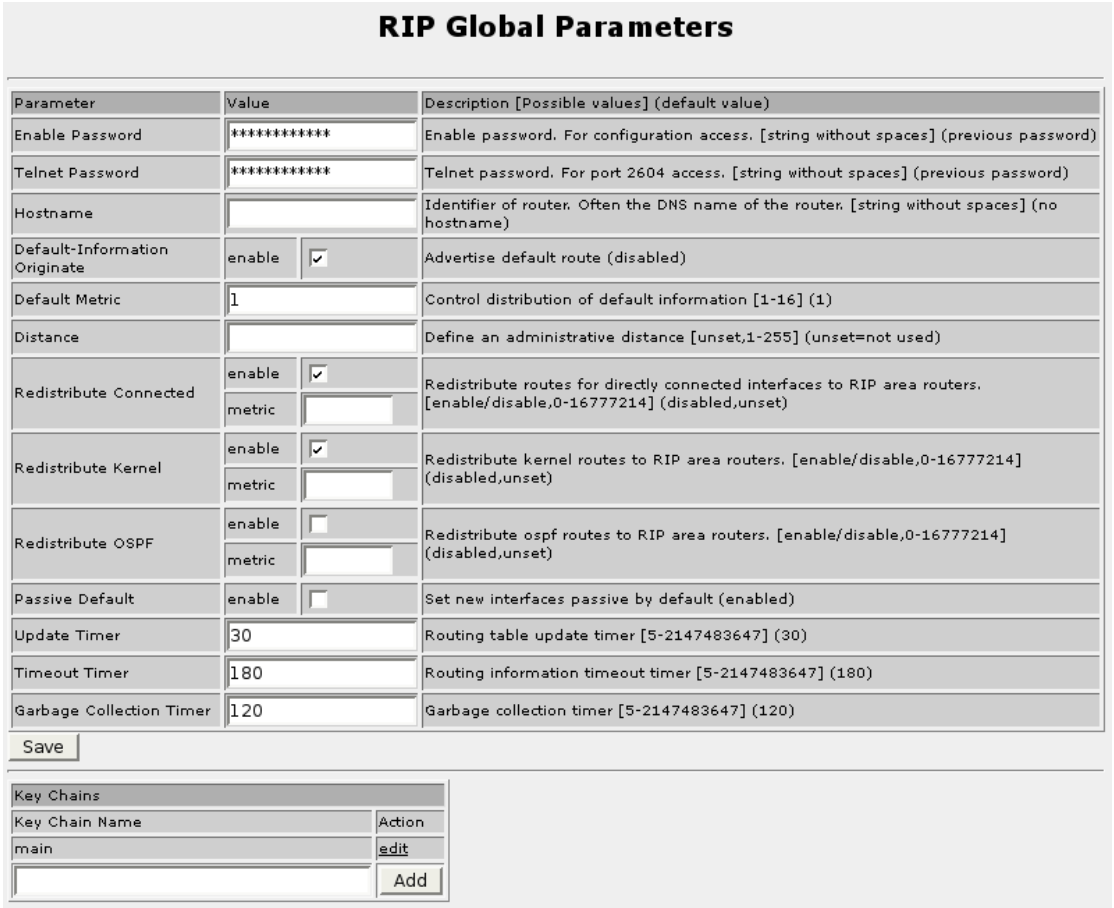

Figure 126: RIP Global Parameters

The **Enable Password** field sets the password to be used for the enable command of ripd. This is used by the telnet interface of ripd to control access to the configuration.

The **Telnet Password** field sets the password to be used for telnet access to ripd. This is used as the login password of ripd when locally telnetting to port 2604 of the router.

The **Hostname** field sets the hostname for the rip daemon. This value is only used as a a reference for convenience. The telnet interface prompt will contain this hostname. The router's system wide hostname is used if this field is left blank.

The **Default-Information Originate** field, when enabled, causes the router to advertise its default route to the RIP network.

The **Default Metric** field sets the default metric to be used for RIP routes which don't have another metric specified.

The **Distance** field sets the administrative distance to use for all routes unless overridden by other distance settings.

The **Redistribute Connected** fields control distribution of connected routes. When enabled, RIP will advertise routes to directly connected interfaces to other RIP routers in the area. Normally only routes that fall within the scope of the network areas will be advertised.

The **Redistribute Kernel** fields control distribution of kernel routes. When enabled, RIP will advertise routes from the kernel routing table, which includes static routes entered by the administrator, to other RIP routers in the area. Normally only routes that fall within the scope of the network areas will be advertised.

The **Redistribute RIP** fields control distribution of routes learned by RIP. When enabled, RIP will advertise routes learned by RIP.

The **Passive Default** option controls the default active/passive state of new interfaces. When enabled all new interfaces will be passive by default. The passive state of individual interfaces is controlled from the RIP Interfaces configuration.

The **Update timer** field controls how often RIP sends out routing table updates.

The **Timeout Timer** field controls how long information stays in the routing table after it is received without an update.

The **Garbage Collection Timer** field controls how long expired entries are remembered before being purged.

RIP Key Chains

The Key Chains table configures authentication keys used on the interfaces. By defining the keys in a key chain, the same settings can be applied to multiple groups of interfaces. Without key chains the same settings would have to be entered for each interface separately.

Key chains also allow multiple keys to be entered in a single key chain with a start time for when that key should become valid as well as the duration the key is valid. This allows multiple keys to be set up with automatic transitions from one key to the next over time.

A key consists of a key string, which is the value used for authentication. It also has the optional lifetime to accept RIP messages with the key, and the optional lifetime to send RIP messages with that key.

RIP Interfaces

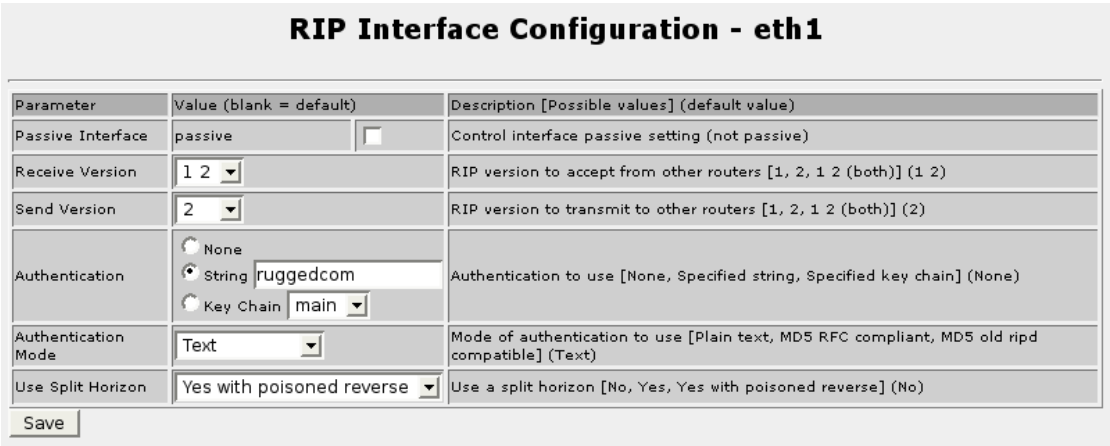

Figure 127: RIP Interfaces

Parameters specific to one interface are configured here.

Each interface on the router is listed. Clicking on settings displays a menu of configuration options for that interface.

Clicking “Remove inactive interfaces” purges the list of any interfaces which are no longer configured on the router.

The **Passive Interface** option controls if an interface is active or passive. Passive interfaces do not send RIP updates to other routers.

The **Receive Version** field controls which versions of RIP messages will be accepted from. Either version 1, 2 or both versions can be accepted. By default both RIP versions are accepted.

The **Send Version** field controls which versions of RIP messages to send to other routers. Either version 1, 2 or both versions can be sent. By default only RIP version 2 messages are sent.

The **Authentication** fields choose the authentication mode this port uses. A port can either use no authentication, use a specific authentication string (used the same was as the string in a key), or use a specific key chain's settings. By default no authentication is used.

The **Authentication mode** field chooses the mode of authentication used. Options are plain text (the default), MD5 following the RIP authentication RFC, and MD5 using the method used by the old ripd implementation.

The **Use Split Horizon** field controls use of the RIP split-horizon feature (RIP v2 only). It can be disabled or enabled, and if enabled it can optionally enable the poisoned reverse feature. Split horizon controls whether routes learned through an interface should be allowed to be advertised back out that interface. By default RIP advertises all routes it knows about to everyone, which makes it take a very long time for dropped links to age out of the network. The split horizon prevents advertising those routes back out the same interface which helps to control this problem. Some network topologies with rings of routers will still have some issues with aging out dead routes even with split horizon enabled but they will still age out faster. If fast network recovery is desired, use OSPF.

RIP Networks

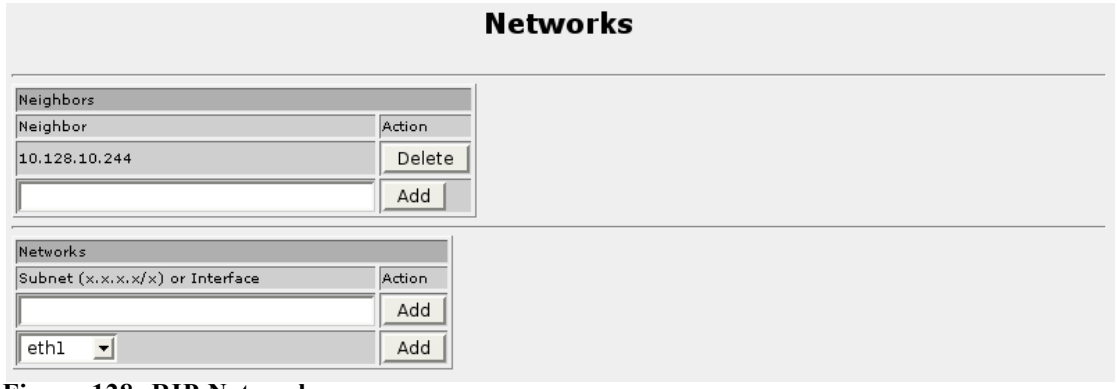

Figure 128: RIP Networks

Neighbors are specific routers with which to exchange routes using the RIP protocol. This can be used when you want to explicitly control which routers are part of your RIP network.

Networks are used when you want to add any router that is part of a specific subnet, or connected to a specific network interface to be part of your RIP network.

Both neighbors and networks can be used at the same time.

Note: *For point to point links (T1/E1 links for example) one must use neighbor entries to add other routers to exchange routes with. Also note that RIP v1 does not send subnet mask information in its updates. Any defined networks are restricted to the classic (in the sense of Class A, B and C) networks. RIP v2 does not have this failing.*

RIP Status

This status menu shows various pieces of information about the current RIP status. The status of each interface is shown, the current database, the current RIP neighbors and the current RIP routing table.

View RIP Configuration

This menu shows the current configuration file of RIP.

This page intentionally blank

Chapter 14 – Configuring Link Backup

Introduction

This chapter familiarizes the user with:

- Configuring link backup
- Obtaining system status
- Testing link backup

Link Backup Fundamentals

Link backup provides an easily configured means of raising a backup link upon the failure of a designated main link. The main and backup links can be Ethernet, CDMA or Dial Modem, TE1, DDS, ADSL or T3. The only requirement is that the main link be a "permanent" link raised at boot time.

The feature can back up to multiple remote locations, managing multiple main:backup link relationships. When the backup link is a modem, many "profiles" of dialed numbers can exist (each serving as a distinct backup link).

The feature can back up a permanent high speed WAN link to a permanent low speed WAN link. This is used when OSPF cannot be employed, such as on public links.

The feature can be used to migrate the default route from the main to the backup link.

The time after a main link failure to backup link startup and the time after a main link recovery to backup link stop are configurable.

The status of the system and a method of testing fail over is provided.

Path Failure Discovery

In order to discover the failure of a primary path (here, through Network A) the link backup daemon will both inspect the link status of the main link and send a regular ping to a designated host. In this way, failures of network links within the cloud are discovered. It is essential that the host always respond to the ping. Another option is to configure a dummy address within the router and ping that address.

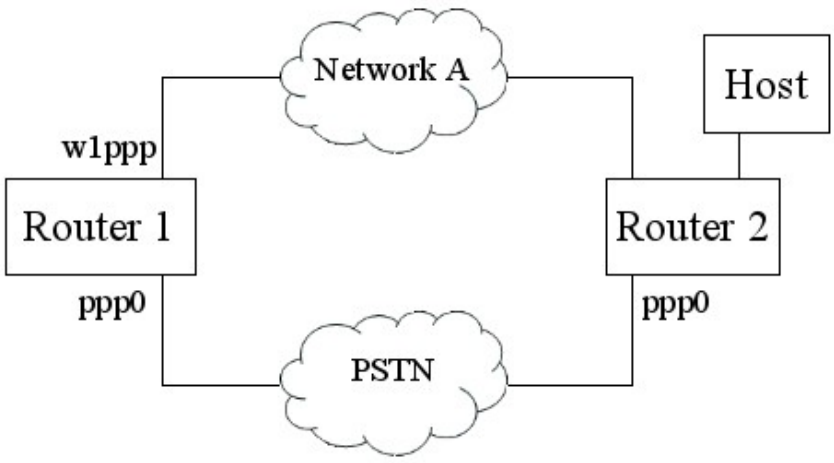

Figure 129: Link Backup Main Menu

The daemon will construe the main link as having failed (even if its link status is "up") if the remote host fails to respond to configurable number of pings after waiting a configurable timeout for each ping.

Use Of Routing Protocols And The Default Route

If the main trunk is on a private network, employ a routing protocol to ensure that an alternate route to end network is learned after the backup trunk is raised. Ensure that OSPF/RIP are configured to operate on the secondary trunk, assigning it a higher metric cost than that of the main trunk.

If the main trunk is on a public network, employ the "transfer default route" feature.

Link Backup Main Menu

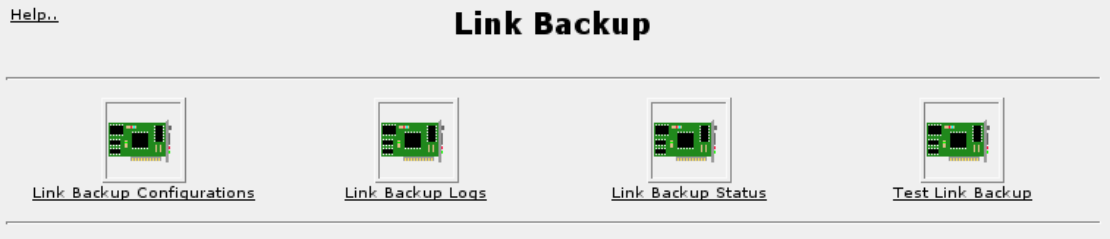

Figure 130: Link Backup Main Menu

Note that Link backup is disabled by default and may be enabled via the System folder, Bootup And Shutdown menu.

Link backup can be configured through the **Link Backup Configuration** link.

Link backup status and logs can be viewed through the **Link Backup Status** and the **Link Backup Log** link after the daemon has been started. A link backup configuration can be tested through the **Link Backup Test** link.

Link Backup Configuration

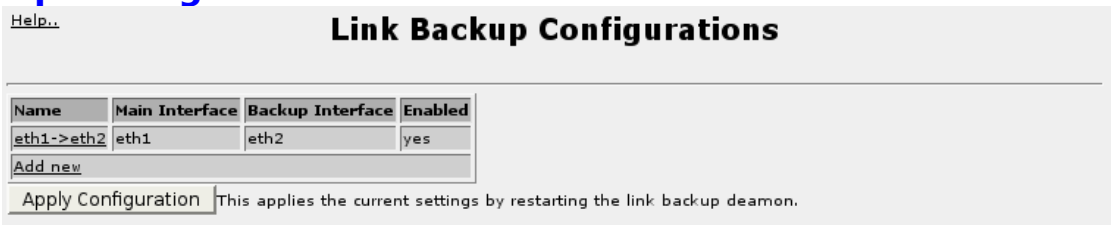

Figure 131: Link Backup Configuration

This menu displays existing main:backup link relationships. Following the links under the **Name** field to an existing pair will edit them or adds a new one.

The **Apply Configuration** button will apply changes by restarting the link backup daemon.

Edit Link Backup Configuration

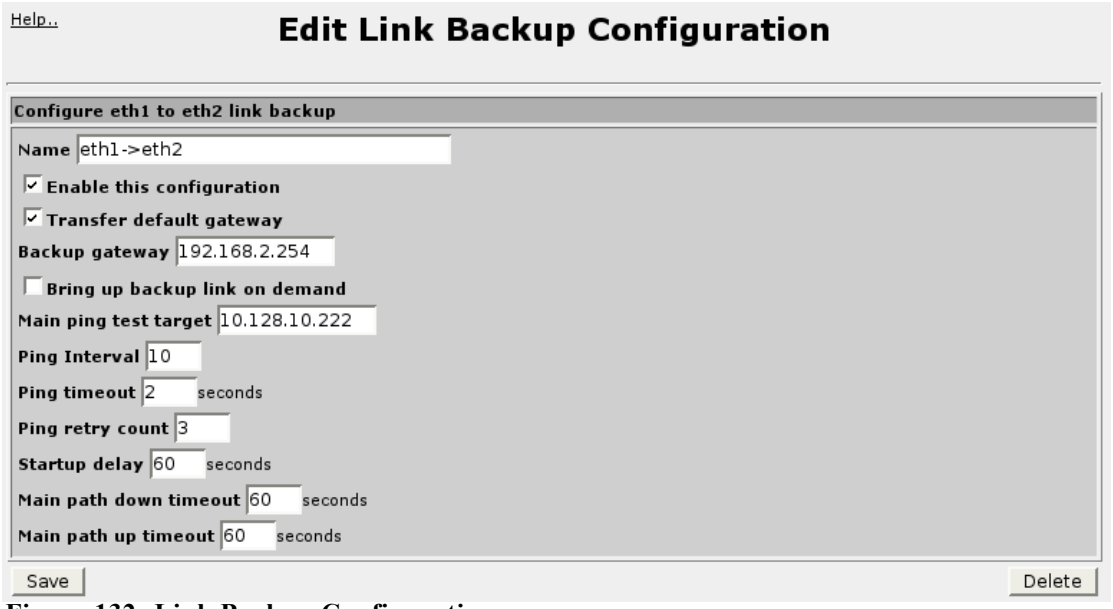

Figure 132: Link Backup Configuration

Set the **Name** field to supply an identification of the pair. This field initially defaults to the "main_link_name->backup_link_name".

The **Enable this configuration** field enables this backup.

The **Transfer default gateway** field causes the gateway to be transferred to the backup link upon failure of the main link path. If the backup interface is point to point, such as PPP, the **Backup gateway** IP address can be automatically determined. Non-point to point interfaces such as Ethernet must be configured with one. The **Bring up backup link on demand** option allows protocols such as DHCP to be used to fetch an address when required.

The **Startup Delay** field configures the length of time to wait for the main link to come up at the start of day.

Note: *If **Startup Delay** is too low, backup will be falsely triggered at start up.*

The **Ping Interval** field configures how often pings are sent.

The **Ping timeout** field configures the duration before immediately retrying a ping.

The **Ping retry count** field configures the number of ping retries before construing a path failure.

Note: *The maximum time to discover a path failure is the length of the **Ping Interval** and the product of the **Number of missed pings before fail over** and the **Ping timeout.***

The **Main path down timeout** field specifies the number of seconds the main trunk must be down before starting the backup trunk.

The **Main path up timeout** field specifies the number of seconds the main trunk must have returned to service before stopping the backup trunk.

You may delete a link backup configuration through the **Delete** button.

Note: *If you delete a link backup configuration that has failed over (or is failing over) to its backup trunk the link daemon will stop attempting the link backup and restore the main trunk, even if the main trunk is still down.*

Link Backup Logs

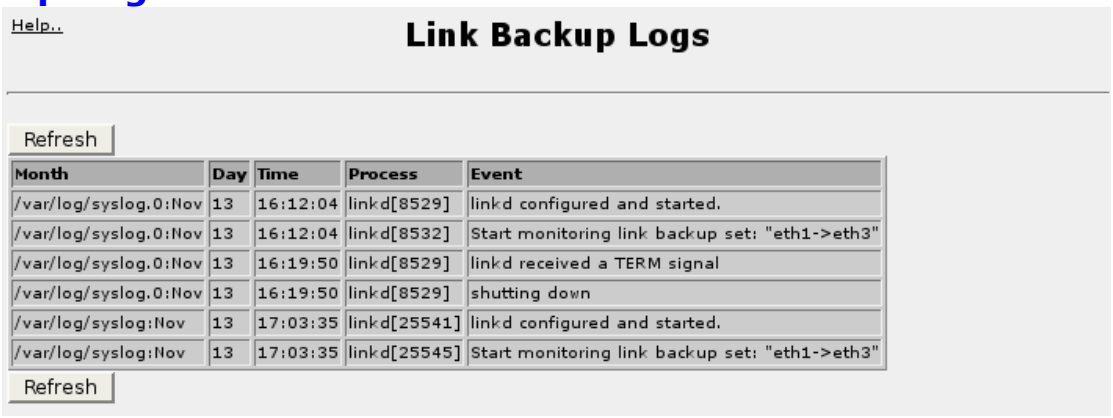

Figure 133: Link Backup Log

The link backup log displays the log of recent backup events.

Link Backup Status

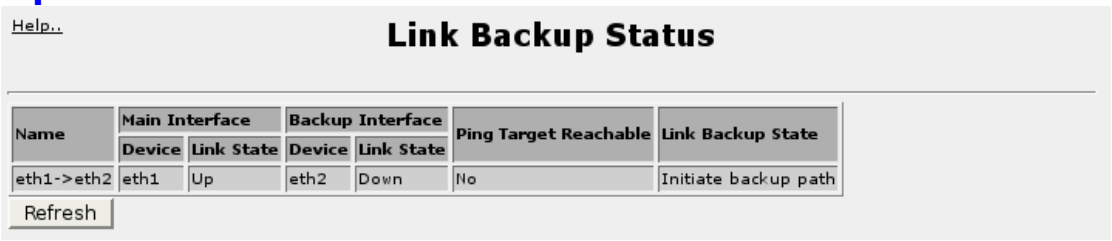

Figure 134: Link Backup Status

The link backup status menu displays the status of links managed by the feature.

Test Link Backup

Figure 135: Test Link Backup

The test link backup menu tests a link backup by discarding all data received on the main interface. This convinces the daemon that the main trunk is unusable and forces it to fail over to the backup trunk.

The **Test Duration** field controls the amount of time to run before restoring service to the main trunk. Please note that this duration must take into account the timing parameters of the backup configuration: The duration should comfortably exceed the **Ping Interval** plus the **Ping Timeout** multiplied by the **Ping retry count** plus the **Main path down timeout**. In the case of a dial backup configuration, also be sure to take into account the call setup and modem connection times. Add to this a time that will allow time to navigate the webmin menus to observe that Link Backup status, link states, and routing are all as expected before, during, and after the Link Backup test.

This page intentionally blank

Chapter 15 – Configuring VRRP

Introduction

This chapter familiarizes the user with:

- Configuring VRRP
- Enabling And Starting VRRP
- Obtaining VRRP Status

VRRP Fundamentals

The Virtual Router Redundancy Protocol (VRRP) eliminates a single point of failure associated with statically routed networks by providing automatic failover using alternate routers. The RuggedRouter VRRP daemon (keepalived) is an RFC 2338 version 2 compliant implementation of VRRP.

The Problem With Static Routing

Many network designs employ a statically configured default route in the network hosts. A static default route is simple to configure, requires little if any overhead to run and is supported by virtually every IP implementation. When dynamic host configuration protocol (DHCP) is employed, hosts may accept configuration for only a single default gateway.

Unfortunately, this approach creates a single point of failure. Loss of the router supplying the default route or the router's WAN connection results in isolating the hosts relying upon the default route.

There are a number of ways that may be used to provide redundant connections to the host. Some hosts can configure alternate gateways while others are intelligent enough to participate in dynamic routing protocols such as Routing Information Protocol (RIP) or Open Shortest Path First routing protocol (OSPF). Even when available, these approaches are not always practical due to administrative and operation overhead.

The VRRP Solution

VRRP solves the problem by allowing the establishment of a “virtual router group”, composed of a number of routers that provide a specific default route. VRRP uses an election protocol to dynamically assign responsibility for the “virtual” router to one of the routers in the group. This router is called the VRRP Master. If the Master (or optionally its WAN connection) fails, the alternate (i.e. backup) routers in the group elect a new Master. The new master provides the virtual IP address and issues a gratuitous ARP to inform the network of where the gateway can be reached.

Because the host's default route does not change and MAC address is updated, packet loss at the hosts is limited to the amount of time required to elect a new router.

VRRP Terminology

Each physical router running VRRP is known as a VRRP Router. Two or more VRRP Routers can be configured to form a “Virtual Router”. Each VRRP Router may participate in one or more Virtual Routers.

Each Virtual Router has a user-configured **Virtual Router Identifier** (VRID) and an **Virtual IP address** or set of IP addresses on the shared LAN. Hosts on the shared LAN are configured to use these addresses as the default gateway.

One router in the Virtual Router Group will be elected as the **Master**, all other routers in the group will be **Backups**.

Each router in the group will run at a specific **Priority**. The router with the highest priority is elected Master. The value of Priority varies from 1 to 255.

VRRP can also monitor a specified interface and give up control of a VRIP if that interface goes down.

In the following network, host 1 uses a gateway of 1.1.1.253 and host 2 uses a gateway of 1.1.1.252. The 1.1.1.253 gateway is provided by VRID 10. In normal practice router 1 will provide this virtual IP as its priority for VRID 10 is higher than that of router 2. If router 1 becomes inoperative or if its w1ppp link fails, it will relinquish control of VRIP 1.1.1.253 to router 2.

In a similar fashion host 2 can use the VRID 11 gateway address of 1.1.1.252 which will normally be supplied by router 2.

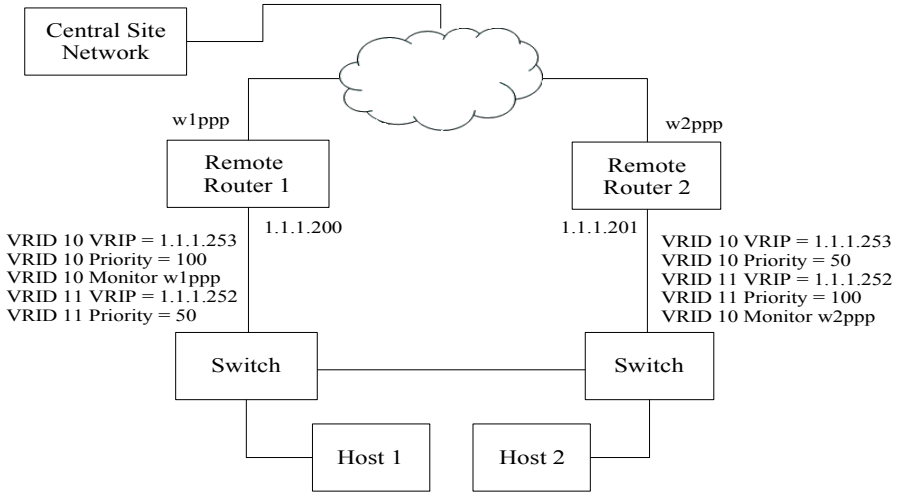

Figure 136: VRRP Example

In this example traffic from host1 will be sent through router 1 and traffic from host2 through router 2. A failure of either router (or its wan link) will be recovered by the other router.

Note that both routers can always be reached by the hosts at their "real" IP addresses.

Other VRRP parameters are the **Advertisement Interval** and **Gratuitous ARP** Delay.

The advertisement interval is the time between which advertisements are sent. A backup router will assume mastership **four advertisement intervals** after the master fails, so the minimum fail-over time is four seconds. If a monitored interface goes down, a master router will immediately signal an election and allow a backup router to assume mastership.

The router issues a set of gratuitous ARPs when moving between master and backup state. These unsolicited ARPs teach the hosts and switches in the network of the current MAC address and port associated with the VRIP. The router will issue a second set of ARPs after the time specified by the Gratuitous ARP delay.

VRRP Main Menu

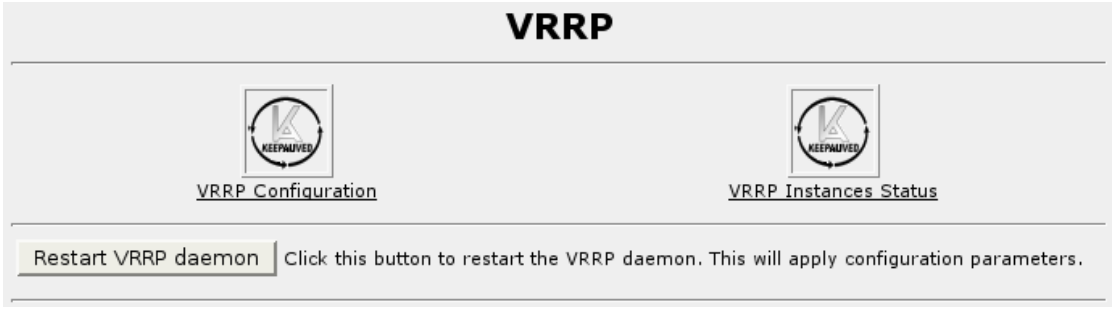

Figure 137: VRRP Main Menu

Note that VRRP is disabled by default and may be enabled via the System folder, Bootup And Shutdown menu.

VRRP can be configured through the **VRRP Configuration** link before the daemon is started.

When enabled, any configuration changes may be made to take effect by selecting the **Restart VRRP daemon** button.

The **VRRP Instances Status** link presents the status VRRP instances existing as of the last restart of keepalived.

VRRP Configuration

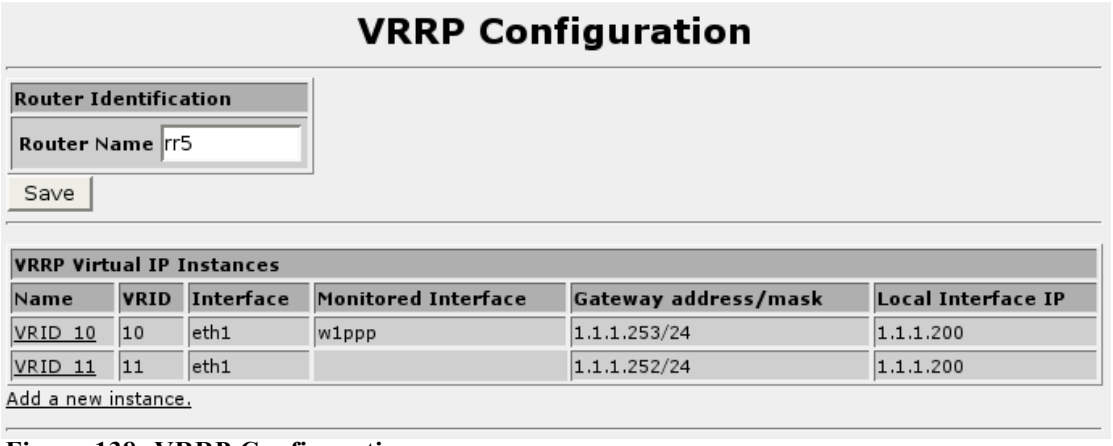

Figure 138: VRRP Configuration

Set the **Router Name** field to supply an identification of the router for VRRP logs. This field initially defaults to the current hostname.

The VRRP instances under the **Name** column define virtual IP groups. Clicking on a link will allow you to edit that instance.

Editing A VRRP Instance

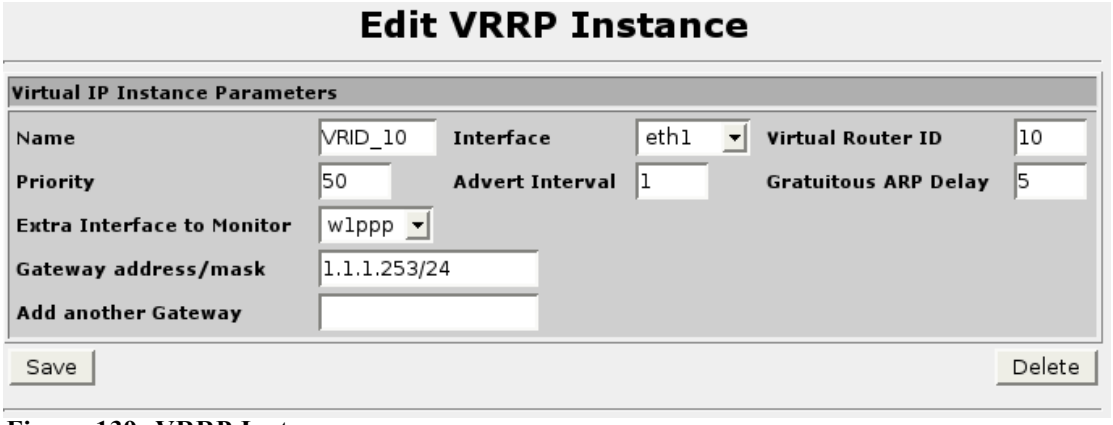

Figure 139: VRRP Instance

The **Name** field is purely for informational purposes.

The **Interface** field configures the interface that VRRP packets are sent upon.

The **Virtual Router ID** field determines the VRID number. Ensure that all routers supplying the same VRIP have the same VRID. The value of the VRID varies from 1 to 255.

The **Advert Interval** field configures the time between VRRP advertisements. Ensure that all routers supplying the same VRID have the same interval.

Note: *VRRP will not work properly if the advertisement interval is different in the master and backup routers.*

The **Gratuitous ARP Delay** field controls the number of seconds after the router changes between master and backup state that a second set of gratuitous ARPs are sent. This mechanism offers a second chance to teach the switching fabric and hosts of the new provider of a gateway address.

The **Extra Interface To Monitor** field causes VRRP to release control of the VRIP if this interface stops running. This prevents the situation where a host forwards information to a gateway router that itself has no way to forward the traffic.

The **Gateway address/mask** and **Add another Gateway** fields configure the VRIP gateway addresses associated with this VRID. Both an IP address and appropriate subnet mask must be provided for each gateway.

The **Save** button saves the virtual instance.

The **Delete** button deletes the virtual instance. After you save or delete an instance you must restart the daemon to action your change.

Viewing VRRP Instances Status

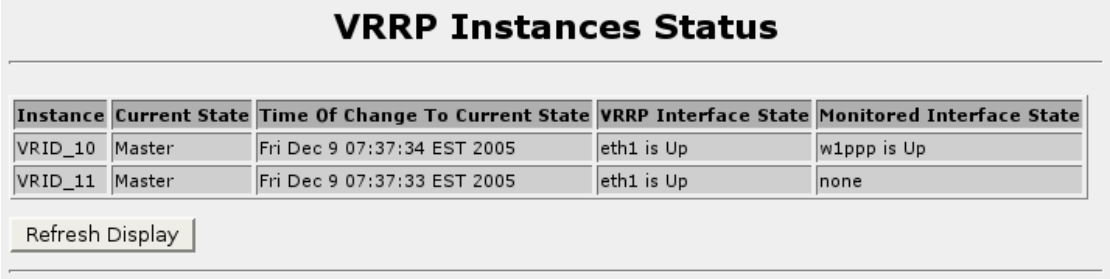

Figure 140: VRRP Instances Status

The VRRP Instances Status menu displays the current status of VRRP instances. This menu does not update status in real time. Click on the **Refresh Display** button to update to the current status.

The entries under the **Instance** column reflect the name of VRRP instances existing as of the last restart of keepalived.

The entries under the **Current State** column reflect the state VRRP instances. An instance can be in one of Master (master for the VRIP), Backup (backup for the VRIP) or Fault (VRRP interface or Monitored interface) is down.

The entries under the **Time Of Change To Current State** column reflect when the current state was entered.

The entries under the **VRRP Interface State** column reflect the link state of the interface that the instance runs upon.

The entries under the **Monitored Interface State** column reflect the link state of the monitored interface or “none” if an interface is not configured.

Chapter 16 – Configuring Traffic Prioritization

Introduction

This chapter familiarizes the user with:

- Enabling/Disabling Traffic Prioritization
- Viewing Traffic Prioritization Statistics

Traffic Prioritization Fundamentals

The RuggedRouter is able to prioritize traffic transmitted on network interfaces (including Ethernet, T1E1, DSL and PPP ports), giving preferential treatment to certain classes of traffic.

It is important to note that prioritization can only be applied to outbound traffic, inbound traffic can not be prioritized.

The two key elements of prioritization are traffic queues and filters. Each prioritized interface has its own unique set of these elements.

Priority Queues

Prioritization establishes a number of queues, each holding packets of differing priority. When the interface is ready to transmit a packet it selects a packet from the highest priority queue first.

If the interface is busy transmitting when packets arrive, they are enqueued in the appropriate queue.

If the interface is not transmitting when the frame arrives to be enqueued, the frame is immediately transmitted. Prioritization will not add additional delay to a stream of packets of differing priority. Prioritization will simply reorder the sequence of transmission of packets to send higher priority packets first.

Note that it is possible in indefinitely stall the transmission of packets from a lower priority queue if a traffic from a higher queue saturates the interface.

Note: *The router mandates that you must have at least a low, normal and high priority queue. Additionally, the high queue must be of higher priority than the normal queue, which must be of higher priority than the low queue.*

Filters

For each packet to be transmitted on a prioritized interface, the packet is compared against each of the filters on that interface until a match is found. The matching filter directs the packet onto a specific queue. If no matching filter is found the packets Type of Service (TOS) bits in its IP header are examined and used.

It is possible to match on source and destination IP address/mask pairs, source and destination port numbers and protocols.

The 0.0.0.0/0 address/mask matches any IP address.

Protocols that can be matched upon include tcp, udp, icmp, ospf, vrrp and ipsec.

TOS Prioritization

The priority of an IP packet can be derived from its Type of Service field. The TOS field has the following format:

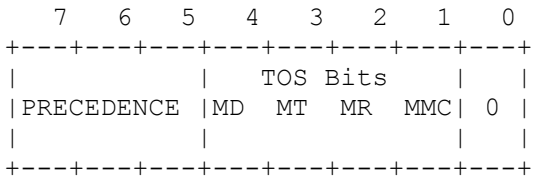

The four TOS bits (the 'TOS field') are defined as:

- MD - Minimize Delay,
- MT - Minimize Throughput,
- MR - Maximize Reliability,
- MMC - Minimize Monetary Cost

As any (or all) of these bits may be set in a packet at a time, there are 16 possible combinations. The router maps these combinations into the high, normal and low priority queues as shown in the following table:

MD	***MT***	***MR***	***MMC***	***Descriptions***	***Priority Queue***
0	0	0	0	Normal Service	Normal
0	0	0	1	Minimize Monetary Cost	Low
0	0	1	0	Maximize Reliability	Normal
0	0	1	1	MR+MMC	Normal
0	1	0	0	Maximize Throughput	Low
0	1	0	1	MT+MMC	Low
0	1	1	0	MT+MR	Low
0	1	1	1	MT+MR+MMC	Low
1	0	0	0	Minimize Delay	High
1	0	0	1	MD+MMC	High
1	0	1	0	MD+MR	High
1	0	1	1	MD+MR+MMC	High
1	1	0	0	MD+MT	Normal
1	1	0	1	MD+MT+MMC	Normal
1	1	1	0	MD+MT+MR	Normal
1	1	1	1	MD+MT+MR+MMC	Normal

Included With Traffic Prioritization

Your RuggedRouter software includes the priostats command line utility, which can be used to show cummulative and one second interval statistics in a format similar to those of the GUI.

Prioritization Example

A remote site router connects to a private network via a T1 line. The router uses OSPF to manage an alternate routing, but its primary purpose is to allow access to a switched network of RuggedServers implementing TcpModbus gateways (TCP/UDP port 502). The router and switches are managed through their Web interfaces, but can me managed through SSH as well. The RuggedServers are managed through Telnet. An SNMP network management polling application tracks the status of all devices.

It is generally wise to ensure that control and management capabilities are always provided. OSPF and SSH/Telnet should be assigned to the highest priority queue. OSPF packet are small and do not consume much bandwidth. SSH and Telnet are not often used but must be available when required.

TcpModbus traffic is ensured a low latency by assigning it the next lowest queue.

Web traffic will be used to manage the router and switches and should be assigned to a still lower queue.

All other traffic can be assigned to a final queue.

In all, four queues are required. The system provides three basic queues (“high”, “normal” and “low”) and a fourth, the “extra high” can be manually added.

Traffic filters are inspected in the order in which they are entered. To reduce load and improve performance the filters should be entered in an order which recognizes the most frequent traffic (under normal conditions). The best filter order is probably:

- match source port 502 -> queue “high”
- match protocol OSPF -> queue “extra high”
- match source port “snmp” -> queue “extra high”
- match source port “www” -> queue “normal”
- match source port “10000” -> queue “normal”
- match source port “ssh” -> queue “extra high”
- match source port “telnet” -> queue “extra high”
- match source IP/Mask 0.0.0.0/0 -> queue “low”

Note that the snmp, www, ssh and telnet keywords are defined in the file /etc/services, so we can use their mnemonics here. We could also have used the raw port numbers 161, 80, 22 and 23, respectively. The TcpModbus port number is not common, and must be explicitly entered. The webmin port number of 10000 reflects the fact that web traffic from a router is issued on this port.

Each of the “port based” filters must match a source port. Matching is being applied to packets from the service at the well known source port to an unknown and variable destination port number.

Finally, note that the final traffic filter essentially suppresses TOS inspection by directing all unmatched traffic onto the “low” queue.

Traffic Prioritization Main Menu

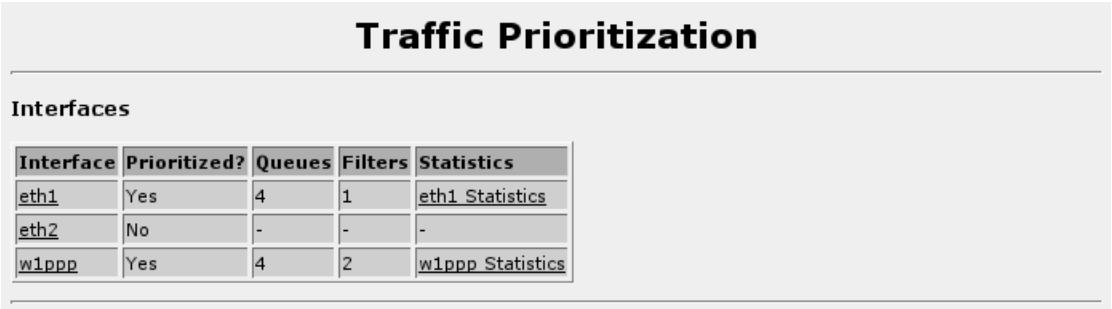

Figure 141: Traffic Prioritization Main Menu

This menu displays network interfaces for which prioritization may be activated. Prioritization may be configured by following the **Interface** column link. The statistics of prioritized interfaces may be viewed by following the links in the **Statistics** column.

Interface Prioritization Menu

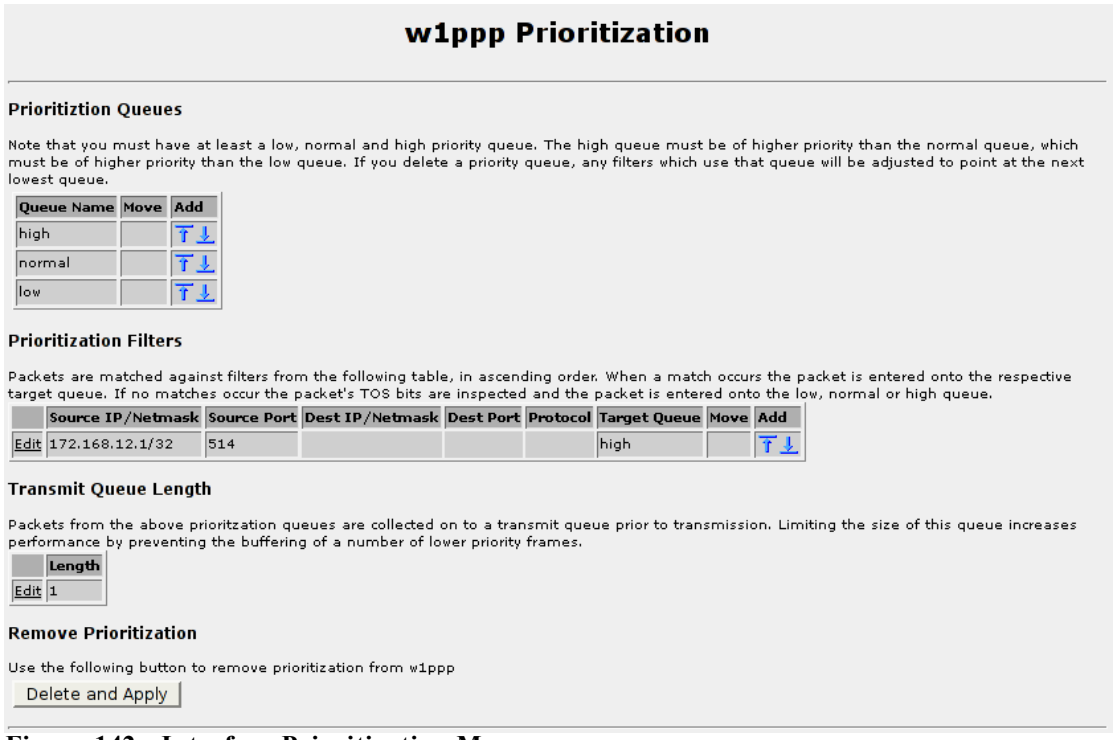

Figure 142: Interface Prioritization Menu

This menu allows you to add, delete and configure queues and filters. Add a new queue or filter by by clicking on the add-above or add-below arrows in the **Add** field. You may also edit a manually created queue by following its link under the **Queue Name** column, and edit a filter by following it's "Edit" link.

Reorder the queues and filters by clicking on the arrows in the **Move** field. Some restrictions apply with queues. You are not allowed to reorder queues in a way that violates the priority implicit in their name.

The Transmit Queue Length Selector allows you to make a tradeoff between latency and performance.

Remove prioritization by selecting the **Delete and Apply** button.

Prioritization Queues

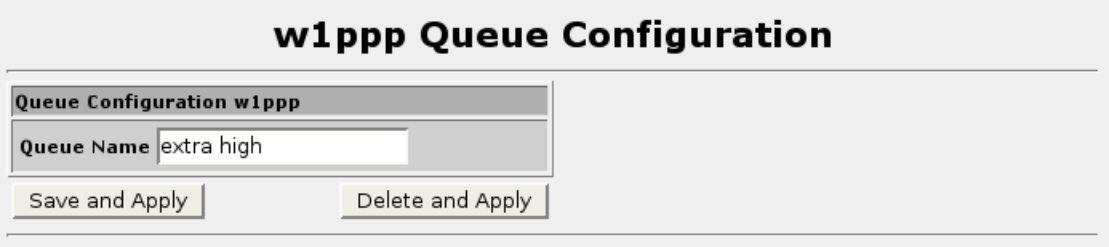

Figure 143: Prioritization Queue Configuration

This menu allows you to edit the name of of a priority queue and to delete the queue.

If you delete a queue referenced by filters, the filters will be adjusted to use the next lowest queue.

Prioritization Filters

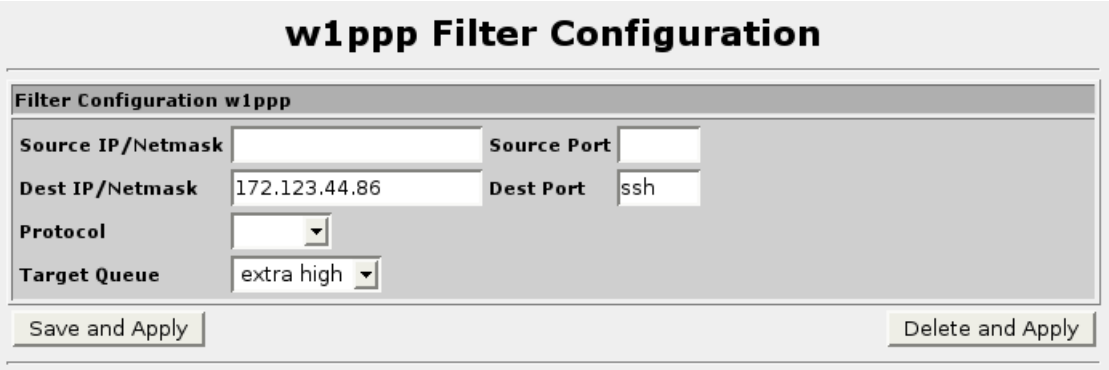

Figure 144: Prioritization Filter Configuration

This menu allows you to edit and delete traffic filters.

The **Source IP/Netmask** and **Dest IP/Netmask** fields specify the IP addresses and masks used to match an outgoing packet. Use 0.0.0.0/0 to generate an “all packets” match.

The **Source Port** and **Dest Port** fields specify the port numbers used to match an outgoing packet. You may specify either a raw number or a mnemonic as specified in the /etc/services file. This setting matches both udp and tcp ports, unless the **Protocol** field specifies udp or tcp.

The **Protocol** field specifies a protocol to match against, currently either tcp, udp, icmp, ospf, vrrp or ipsec.

The **Target Queue** field selects one of the available priority queues.

Prioritization Transmit Queue Length

The WAN protocols supplied by the RuggedRouter rely upon transmit queues to ensure their efficiency. Even as a packet is starting to be transmitted, other packets can be lining up behind it. Normally there is only one queue, the transmit queue, and packets are transmitted from it in the order in which they arrived.

The transmit queue is a means of enhancing performance.

Prioritization favors some packets over others by transmitting them with preference.

Prioritization works by establishing queues at the required priority levels filling the transmit queue with them in priority order. The aim of establishing low latency for certain traffic is foiled when transmit queue lengths are large because multiple low priority packets may have queued before a high priority packet arrives at the router.

RuggedCom recommends that the transmit queue length be left at its minimum default value of 1. Higher values, however, may strike a balance between latency an performance.

Prioritization Statistics

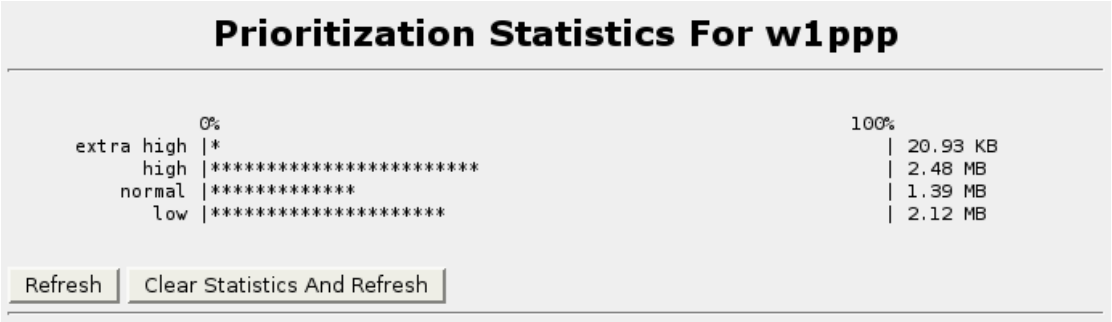

Figure 145: Prioritization Statistics

This menu displays the percentage of interface traffic that has been transmitted from each priority queue. The **Refresh** button causes the statistics to be updated. The **Clear Statistics And Refresh** button causes the statistics to be cleared and then captured after a one second interval.

Chapter 17 – Configuring Generic Routing Encapsulation

Introduction

This chapter familiarizes the user with:

- Enabling/Disabling GRE
- Viewing GRE Status

GRE Fundamentals

The RuggedRouter is able to encapsulate multicast traffic and IPv6 packets and transport them through an IPv4 network tunnel.

The GRE tunnel can transport the traffic through any number of intermediate networks. The key parameters for GRE in each router are the tunnel name, local router address, remote router address and remote subnet.

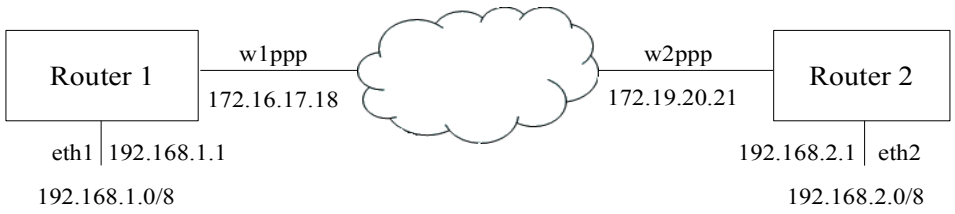

Figure 146: VRRP Example

In the above example, Router 1 will use a GRE tunnel with a local router address of 172.16.17.18, a remote router address of 172.19.20.21, and a remote subnet of 192.168.2.0/24.

Note: *If you are connecting to a CISCO router, the local router address corresponds to the CISCO IOS "source" address and the remote router address corresponds to the "destination" address.*

You may also set a cost for the tunnel. If another method of routing between Router1 and Router2 becomes available, the tunneled packets will flow through the lowest cost route. You can optionally restrict the packets by specifying the local egress device (in the case of router1, w1ppp).

GRE Main Menu

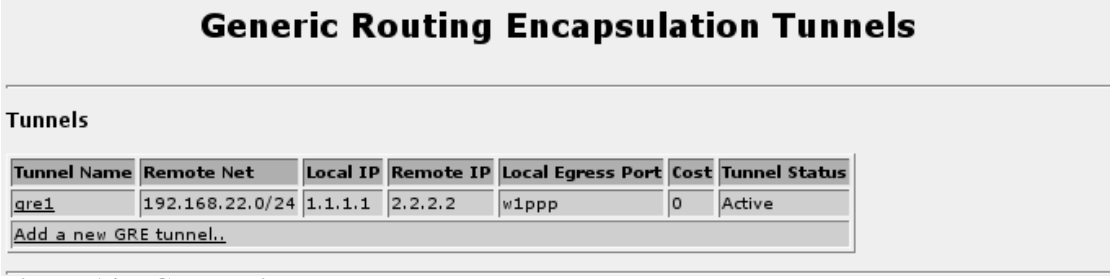

Figure 147: GRE Main Menu

This menu displays configured GRE tunnels. The tunnel status will be "active" if the tunnel was successfully created.

GRE Configuration Menu

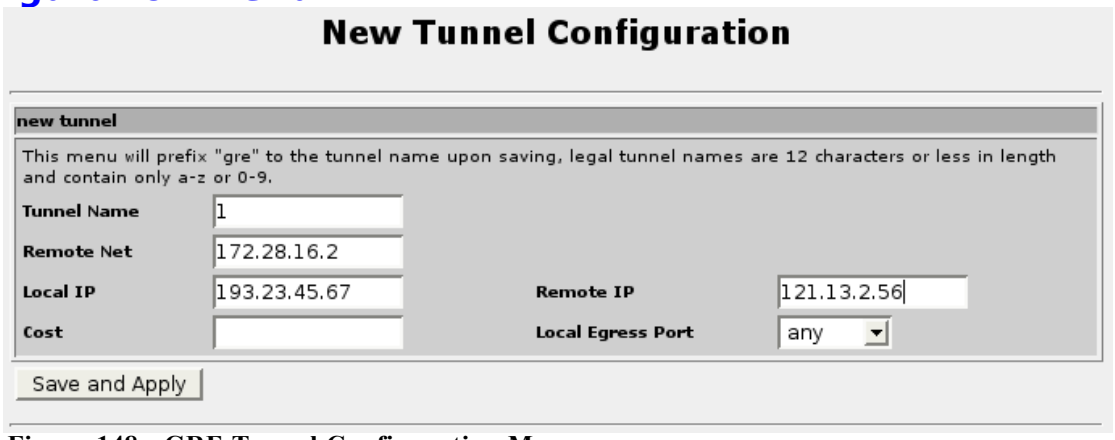

Figure 148: GRE Tunnel Configuration Menu

This menu allows you to add or edit a tunnel.

The **Tunnel Name** field will be presented if the tunnel is being created. The tunnel name is purely for informational purposes. A network routing device with this name will be created. In order that the name not collide with those used by other interfaces, it will be prefixed with "gre".

The **Remote Net** field configures the target network whose traffic is forward through the tunnel. It may be a individual IP address or subnetted IP address such as 192.168.0.0/24. The Remote Net must not used by another tunnel.

The **Local IP** field configures the IP address of the local end of the tunnel.

The **Remote IP** field configures the IP address of the local remote of the tunnel.

Note: *Each tunnel must have a unique combination of local and remote addresses, or it will not be activated.*

The **Cost** field configures the routing cost associated with networking routing that directs traffic through the tunnel. The cost will default to zero if left unset.

The **Local Egress Port** configures a port to bind the tunnel to. If set, tunneled packets will only be routed via this port and will not be able to escape to another device when the route the to endpoint changes.

This page intentionally blank

Chapter 18 – Network Utilities

Introduction

This chapter familiarizes the user with:

- Pinging hosts,
- Running a traceroute,
- Performing a host lookup,
- Tracing line activity,
- Showing interface statistics.

Network Utilities Main Menu

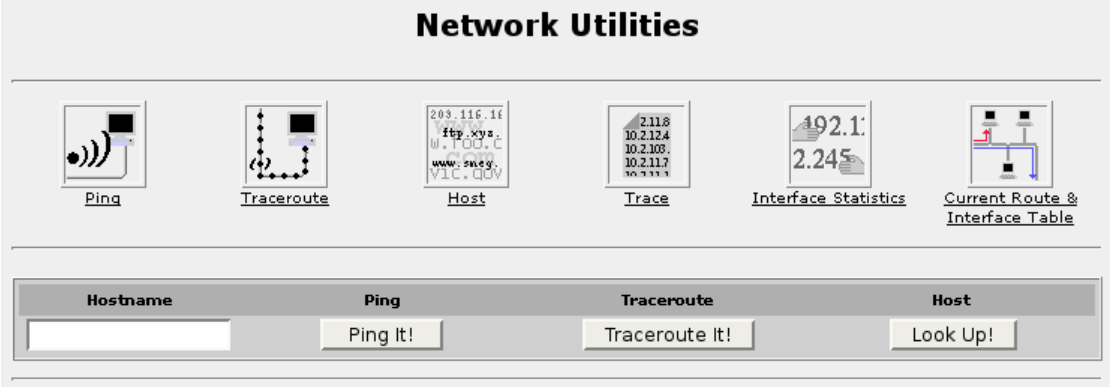

Figure 149: Network Utilities Main Menu

The lower part of the menu provides quick pinging, tracerouting and lookup of hosts.

The upper part leads to menus providing more configurable options for these commands. Additionally, Ethernet, WAN and Serial port tracing is provided. A summary of interface statistics and the current routing table is provided.

Ping Menu

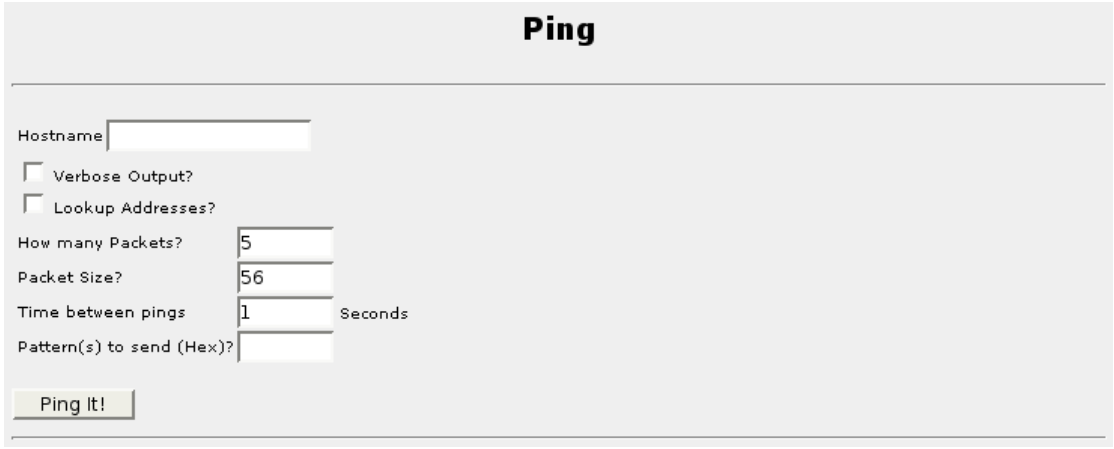

Figure 150: Ping Menu

The **Hostname** field accepts the host name or IP address to ping.

The **Verbose Output?** field causes ping to present the maximum of output.

The **Lookup Addresses?** field causes ping to resolve IP addresses to domain names. This can make ping behave very slowly if DNS is not properly configured.

The **Packet Size?** field specifies the size of the data in the ping packet. The true length of the packet is 28 bytes larger due to IP/ICMP overhead.

The **Time between pings** field limits the rate at which pings are sent.

The **Pattern(s) to send (Hex)?** field specifies a pattern to fill the packet sent. This is useful for diagnosing data-dependent problems in a network. For example, specifying "ff" will cause the sent packet to be filled with all ones.

Traceroute Menu

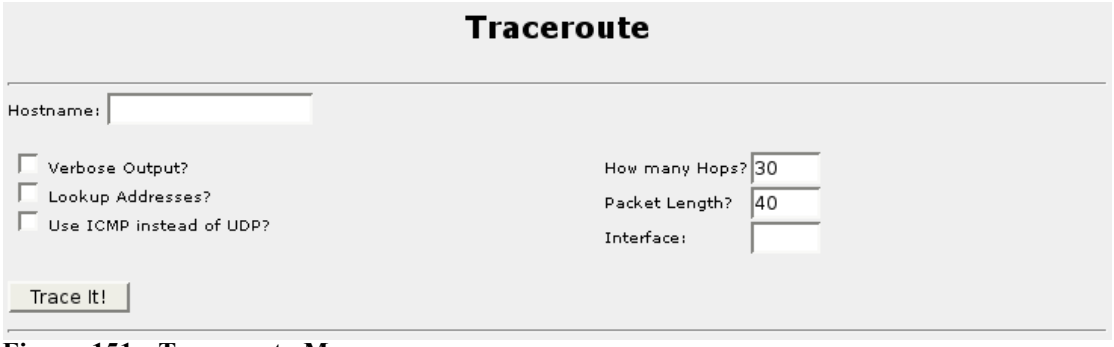

Figure 151: Traceroute Menu

The **Hostname** field accepts the host name or IP address to trace the route to.

The **Verbose Output?** field causes ping to present the maximum of output.

The **Lookup Addresses?** field causes ping to resolve IP addresses to domain names. This can make ping behave very slowly if DNS is not properly configured.

The **Use ICMP instead of UDP?** field causes traceroute to probe with ICMP packets.

The **How many Hops?** field limits the maximum number of hops that traceroute will attempt to map.

The **Packet Length?** field specifies the size of the data in the traceroute packet.

The **Interface?** field specifies the network interface to obtain the source IP address for outgoing probe packets. Otherwise the router will manually set the address based on the actual interface taken.

Host Menu

Figure 152: Host Menu

The **Hostname** field accepts the host name or IP address to ping.

The **Type?** field selects the type of information to capture.

The **Nameserver?** fields select the server to use to resolve with. If **Default** is left selected the DHCP, DNS or local resolv.conf setup will be used. Otherwise the address supplied will be used.

The **Timeout?** field specifies the maximum time to wait before abandoning a lookup.

Trace Menu

The trace menu contains three sections providing the the capability to trace network interfaces, Frame Relay Interfaces and Serial server interfaces. The latter two menus will appear only if you have configured Frame Relay or Serial server interfaces.

Tcpdump A Network Interface

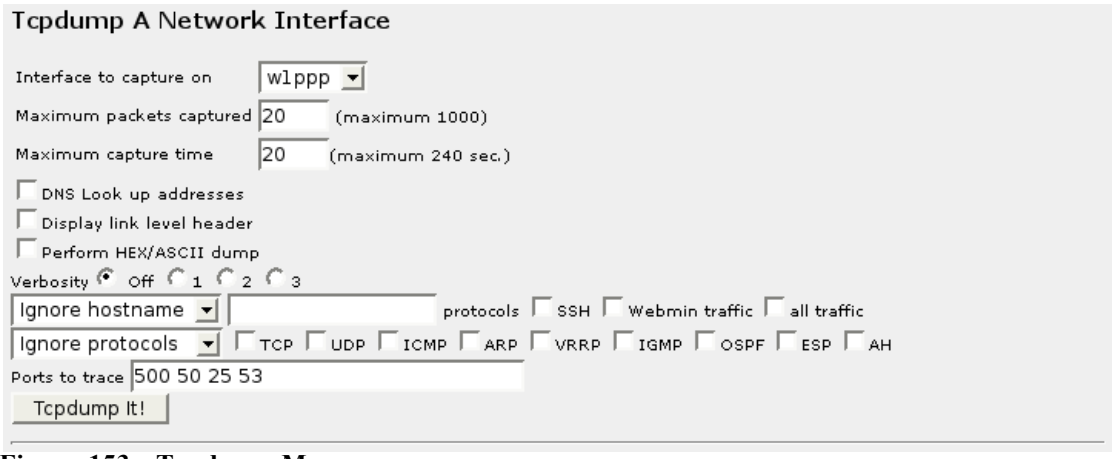

Figure 153: Tcpdump Menu

The **Interface to capture on** field specifies the interface to show traffic on.

The **Maximum packets captured** and **Maximum capture time** fields limits the amount of traffic captured.

The **Lookup Addresses?** field causes ping to resolve IP addresses to domain names. This can make ping behave very slowly if DNS is not properly configured.

The **Display link level header** field causes this header to be displyed.

The **Perform HEX/ASCII dump** field will cause the data content of the captured packets to be displayed. This option generates a large amount of data.

The **Verbosity** fields specify the level of decoding which tcpdump supplies.

The **Ignore hostname/Only hostname** selector excludes or selects the IP address specified in the next field. If the **SSH** box is selected, SSH traffic from The IP will be excluded/shown. If the **Webmin traffic** box is selected, Webmin traffic from The IP will be excluded/shown. If the **All traffic** box is selected, traffic from The IP will be excluded/shown. **This option provides a filter capability to tcpdump an interface and to block the users own traffic from being displayed.**

The **Ignore protocols/Only protocols** selector excludes or selects the protocols specified in the next fields.

The **Ports to trace** field specifies TCP/UDP ports to trace. Enter a list of ports separated by spaces to trace more than a single port.

Frame Relay Link Layer Trace A WAN Interface

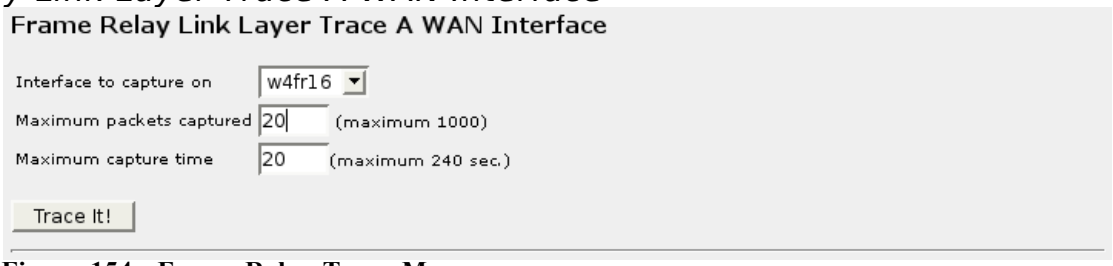

Figure 154: Frame Relay Trace Menu

Frame Relay tracing uses the wanpipemon utility.

The **Interface to capture on** field specifies the interface to show traffic on.

The **Maximum packets captured** and **Maximum capture time** fields limits the amount of traffic captured.

Serial Trace A Serial Server Port

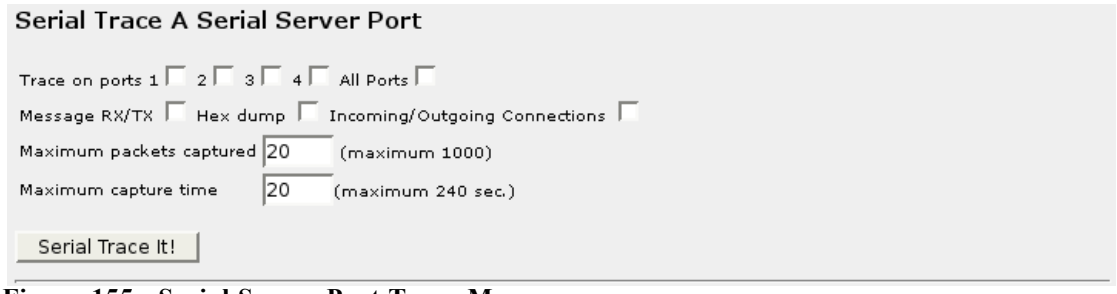

Figure 155: Serial Server Port Trace Menu

The **Trace on ports** fields specify the serial port to show traffic on.

The **Message RX/TX** and **Incoming/Outgoing Connections** fields causes data packets and **Connection activity** to be included in the trace. The **Hex dump** field causes the content of data packets to be displayed.

The **Maximum packets captured** and **Maximum capture time** fields limits the amount of traffic captured.

Interface Statistics Menu

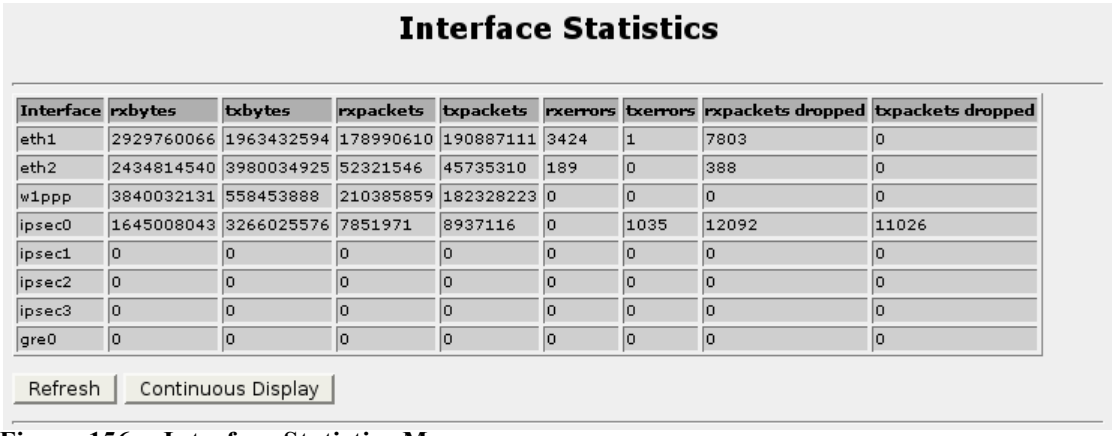

Figure 156: Interface Statistics Menu

This menu provides basic statistics for all network interfaces.

The **Refresh** button will cause the page to be reloaded.

The **Continuous Display** button will cause the browser to continuously reload the page showing the differences in statistics from the last display. **The difference is not a real time rate in bytes or packets per second.**

Note that detailed statistics for T3, T1/E1, DDS and ADSL are available within the menus that configure those interfaces.

Current Routing & Interface Table

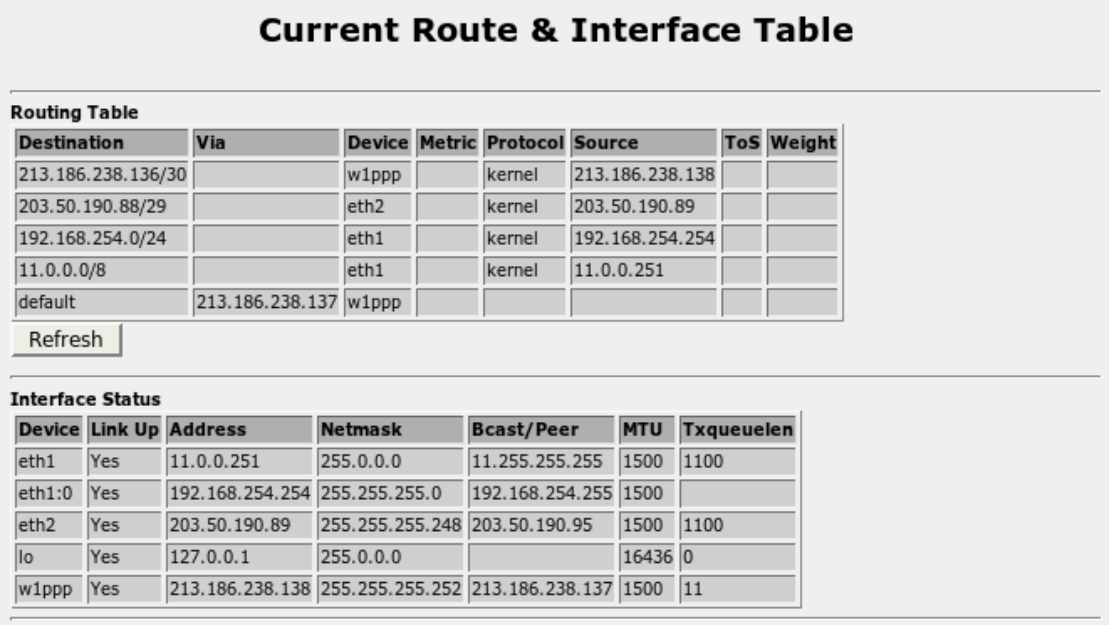

Current Route & Interface Table

Routing Table

Destination	Via	Device	Metric	Protocol	Source	ToS	Weight
213.186.238.136/30		w1ppp		kernel	213.186.238.138		
203.50.190.88/29		eth2		kernel	203.50.190.89		
192.168.254.0/24		eth1		kernel	192.168.254.254		
11.0.0.0/8		eth1		kernel	11.0.0.251		
default	213.186.238.137	w1ppp					

Refresh

Interface Status

Device	Link Up	Address	Netmask	Bcast/Peer	MTU	Txqueuelen
eth1	Yes	11.0.0.251	255.0.0.0	11.255.255.255	1500	1100
eth1:0	Yes	192.168.254.254	255.255.255.0	192.168.254.255	1500	
eth2	Yes	203.50.190.89	255.255.255.248	203.50.190.95	1500	1100
lo	Yes	127.0.0.1	255.0.0.0		16436	0
w1ppp	Yes	213.186.238.138	255.255.255.252	213.186.238.137	1500	11

Figure 157: Current Routing & Interface Table

This menu displays the current routing table and the state of the router's interfaces.

Select the Refresh link in order to refresh the display.

The entries under the **Destination** field reflect the network or host which can be reached through this route. The “default” entry matches any packet which has not already matched another route.

The entries under the **Via** field reflect the address of the gateway to route packets through to reach the target network. The field is blank for non-gateway routings.

The entries under the **Device** field reflect the name of the interface this route belongs to. Packets using this route will be sent on this interface.

The entries under the **Metric** field reflect the the cost of this route. The route with the lowest metric matching a destination is used.

The entries under the **Protocol** field reflect the system that created the route. It is one of “kernel” (default interface routes), “core” (dynamic routing protocol routes), “redirect” (routes added due to ICMP redirect message) or “static” (for manually added routes).

The entries under the **Source** field reflect the source address to use when originating a packet to a destination matching this route. Note that packets routed through the router have their own source address. Note that if the sending application decides to, it can manually specify the source address.

The entries under the **ToS** field reflect the ToS value a packet must match to be routed by this route.

The entries under the **Weight** field reflect the relative bandwidth or quality of this link within a multi-path route. Note that multi-path routes are shown with multiple lines for a single destination.

Interface Status

This menu also summarizes the interface status.

The entries under the **Device** field reflect the name of the device.

The entries under the **Link up** field reflect the current link state of interface.

The entries under the **Address** field reflect the local address of interface.

The entries under the **Netmask** field reflect the netmask applied to this interface.

The entries under the **Bcast/Peer** field reflect the broadcast address for the interface or the peer address if the interface is a point to point interface.

The entries under the **MTU** field reflect the Maximum Transmission Unit size for the interface.

The entries under the **Txqueuelen** field reflect the transmit queue length for the interface.

This page intentionally blank

Chapter 19 – Configuring Serial Protocols

Introduction

This chapter familiarizes the user with:

- RawSockets Applications
- Configuring Serial ports for RawSocket
- Viewing Serial Port and TCP Connection status and statistics
- Resetting Serial ports
- Tracing Serial Port activity

Serial IP Port Features

RuggedCom Serial IP provides you with the following features:

- Raw Socket Protocol -A means to transport streams of characters from one serial port on the router, to a specific remote IP address and TCP port.
- 4 independent serial ports per product
- Baud rates of 1200, 2400, 4800, 9600, 19200, 38400, 57600, 115200 or 230400 bps.
- Supports RS232, RS422 and RS485 party line operation.
- XON/XOFF flow control.
- Support a point-to-point connection mode and a broadcast connection mode in which up to 32 remote servers may connect to a central server.
- TCP/IP incoming, outgoing or both incoming/outgoing connections mode, configurable local and remote TCP port numbers.
- Packetize and send data on a full packet, a specific character or upon a timeout.
- Support a “turnaround” time to enforce minimum times between messages sent out the serial port.
- Debugging facilities include connection tracing and statistics.

LED Designations

The Quad TriplePlay Serial card includes transmit and receive LEDs. The transmit LED is leftmost when the card is in the top slot and will light while characters are being transmitted. The receive LED is rightmost when the card is in the top slot and will light while characters are being received.

Serial port numbers are as described by the “SER” labels as shown in the home page chassis diagram.

Serial Protocols Applications

Character Encapsulation

Character encapsulation is used any time a stream of characters must be reliably transported across a network.

The character streams can be created by any serial device. The baud rates supported at either server need not be the same. If configured, the router will obey XON/XOFF flow control from the end devices.

One of the routers is configured to listen to TCP connection requests on a specific TCP port number. The other server is configured to connect to its peer on the listening port number. The RuggedRouter will attempt to connect periodically if the first attempt fails and after a connection is broken.

The RuggedRouter can be used to connect to any device supporting TCP (e.g. a host computer's TCP stack or a serial application on a host using port redirection software).

RTU Polling

The following applies to a variety of RTU protocols besides ModBus RTU, including ModBus ASCII and DNP.

The remote router communicates with host equipment through:

- native TCP connections,
- another RuggedRouter's via a serial port or
- a port redirection package which Supports TCP.

If a RuggedRouter is used at the host end, it will wait for a request from the host, encapsulate it in a TCP message and send it to the remote side. There, the remote RuggedRouter will forward the original request to the RTU. When the RTU replies the RuggedRouter will forward the encapsulated reply back to the host end.

ModBus does not employ flow-control so XON/XOFF should not be configured.

The RuggedRouter maintains configurable timers to help decide replies and requests are complete and to handle special messages such as broadcasts.

The RuggedRouter will also handle the process of line-turnaround when used with RS485.

Broadcast RTU Polling

Broadcast polling allows a single host connected RuggedRouter to "fan-out" a polling stream to a number of remote RTUs.

The host equipment connects via a serial port to a RuggedRouter. Up to 32 remote RuggedRouters may connect to the host server via the network.

Initially, the remote servers will place connections to the host server. The host server in turn is configured to accept the required number of incoming connections.

The host will sequentially poll each RTU. Each poll received by the host server is forwarded (i.e. broadcast) to all of the remote servers. All RTUs will receive the request and the appropriate RTU will issue a reply. The reply is returned to the host server, where it is forwarded to the host.

Serial Protocols Concepts And Issues

Host And Remote Roles

RuggedRouter either places a TCP connection or accepts one. The connection can be made from the field or "remote" equipment to the central site or "host" equipment, vice versa or bi-directionally.

Connect from the host to the remote if:

- The host end uses a port redirector that must make the connection.
- The host end is only occasionally activated and will make the connection when it becomes active.
- A host end firewall requires the connection to be made outbound.

Connect from the remote to the host if the host end accepts multiple connections from remote ends in order to implement broadcast polling.

Connect from each side to other if both sides support this functionality.

Use Of Port Redirectors

Port redirectors are PC packages that emulate the existence of communications ports. The redirector software creates and makes available these "virtual" COM ports, providing access to the network via a TCP connection.

When a software package uses one of the virtual COM ports, a TCP connection is placed to a remote IP address and TCP port that has been programmed into the redirector. Some redirectors also offer the ability to receive connections.

Message Packetization

The server buffers received characters into packets in order to improve network efficiency and demarcate messages.

The server uses three methods to decide when to packetize and forward the buffered characters to the network:

- Packetize on Specific Character,
- Packetize on timeout and
- Packetize on full packet.

If configured to packetize on a specific character, the server will examine each received character and will packetize and forward upon receiving the specific character. The character is usually a <CR> or an <LF> character but may be any ASCII character.

If configured to packetize on a timeout, the server will wait for a configurable time after receiving a character before packetizing and forwarding. If another character arrives during the waiting interval, the timer is restarted. This method allows characters transmitted as part of an entire message to be forwarded to network in a single packet, when the timer expires after receiving the very last character of the message. This is usually the only packetizer selected when supporting ModBus communications.

Finally, the server will always packetize and forward on a full packet, i.e. when the number of characters fills its communications buffer (1024 bytes).

Use of Turnaround Delays

Some RTU protocols (such as ModBus) use the concept of a turnaround delay. When the host sends a message (such as a broadcast) that does not invoke an RTU response, it waits a turnaround delay time. This delay ensures that the RTU has time to process the broadcast message before it has to receive the next poll.

When polling is performed, network delays may cause the broadcast and next poll to arrive at the remote server at the same time. Configuring a turnaround delay will enforce a minimum separation time between each message sent out the port. Note that turnaround delays do not need to be configured at the host computer side and may be disabled there.

Serial Protocols Main Menu

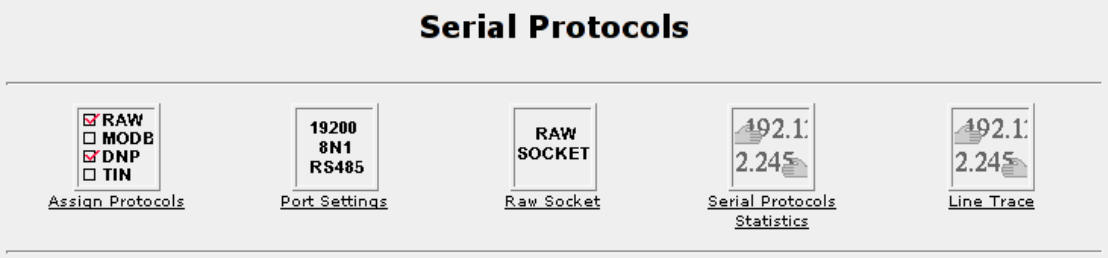

Figure 158: Serial Protocols Server Main Menu

Note that the Serial Protocols server is disabled by default and may be enabled via the System folder, Bootup And Shutdown menu.

The **Assign Protocols** menu assigns a serial protocol to one of your serial ports.

The **Port Settings** menu configures the serial port and its electrical protocol.

If any of your serial ports are configured as **RawSocket** protocol, this menu will configure them.

The **Serial Protocols Statistics** menu will show you the status and statistics for any established sessions.

The **Line Trace** menu will provide a line activity trace for the serial ports.

Assign Protocols Menu

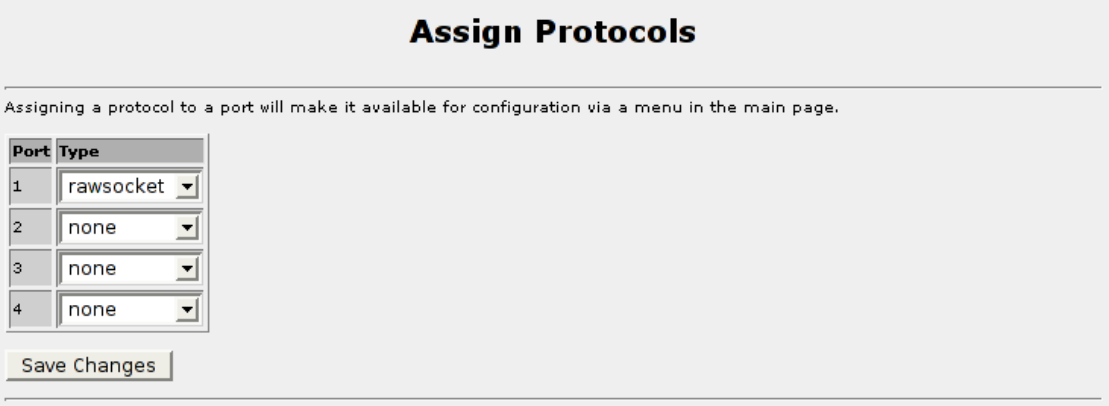

Figure 159: Assign Protocols Menu

This menu associates a protocol with a serial port. Unused ports should be left associated with "none". Changing an association will immediately close the calls of the old protocol.

Port Settings Menu

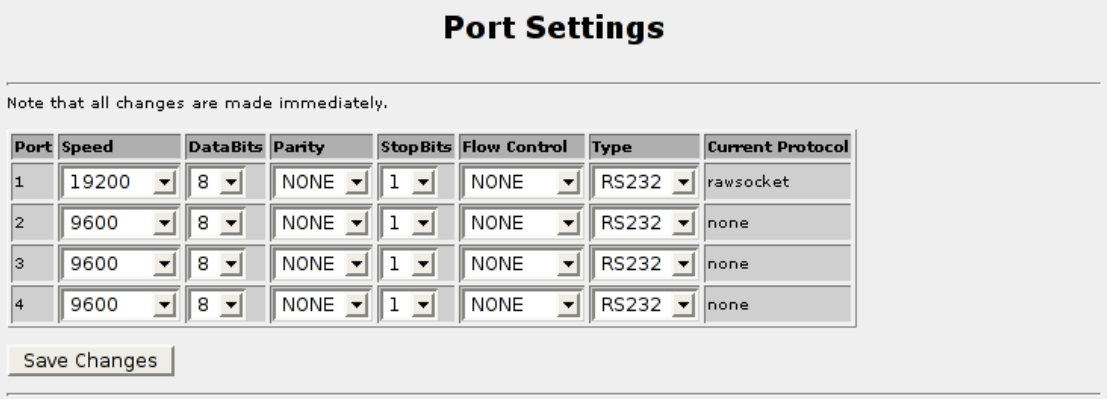

Figure 160: Port Settings Menu

This menu configures the serial settings and electrical protocol associated with a serial port. Changes are made immediately.

RawSocket Menu

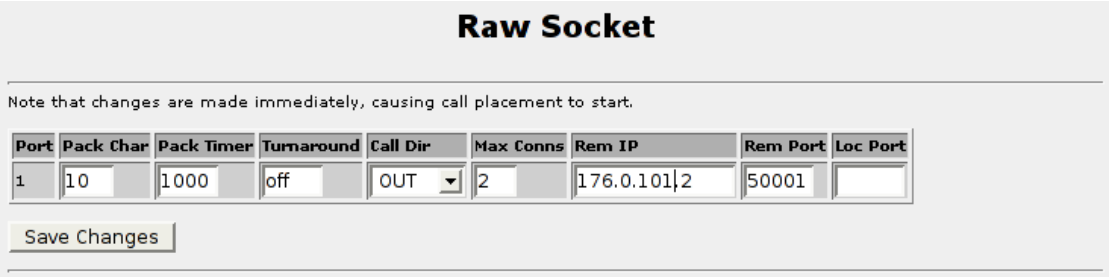

Figure 161: Raw Socket Menu

This menu configures the Raw Socket settings for each port. Changes are made immediately.

The **Pack Char** field configures the numeric value of the ASCII character which will force forwarding of accumulated data to the network. The Pack Char must be between 0 and 255 inclusive or the value off. If configured off, accumulated data will be forwarded based upon the packetization timeout parameter.

The **Pack Timer** field configures the delay from the last received character until when data is forwarded. The Pack Timer must be between 5 and 1000 milliseconds inclusive.

The **Turnaround** timer field controls the amount of delay (if any) to insert between the transmissions of individual messages out the serial port. The Pack Timer must be between 1 and 1000 milliseconds inclusive, of off.

The **Call Dir** field configures whether to accept an incoming connection, place an outgoing connection or do both.

The **Max Conns** field configures the maximum number of incoming connections to permit when the call direction is incoming.

The **Remote IP** field configures the address used when placing an outgoing connection.

The **Remote Port** field selects the TCP destination port used in outgoing connections.

The **Local Port** field selects the local TCP port to use to accept incoming connections.

Serial Protocols Statistics Menu

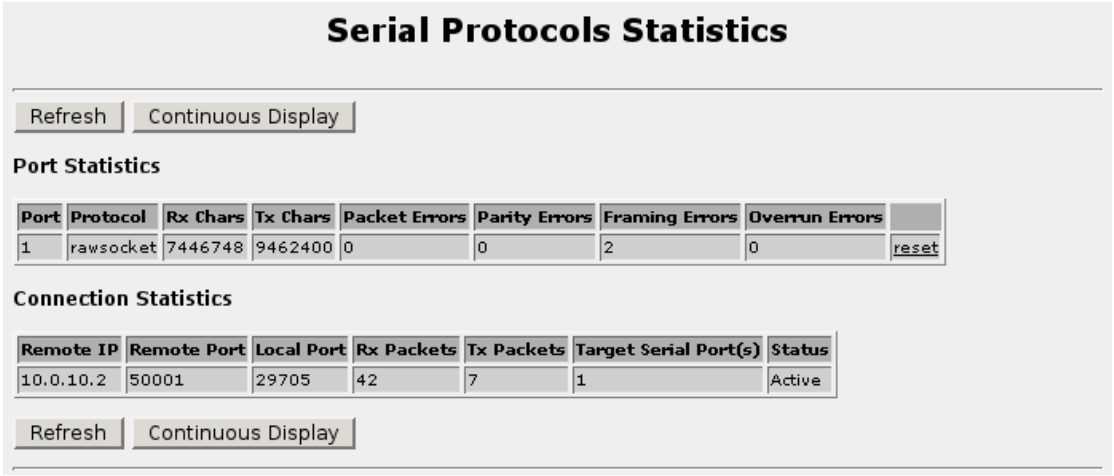

Figure 162: Serial Protocols Statistics Menu

This menu presents statistics of serial port activity and established connections. The menu also allows you to reset a port, forcing call hang-up and re-establishment.

The **Port Statistics** table provides a record for each active serial port. The number of historical received and transmitted characters as well as errors will be displayed.

The **Connection Statistics** table reflects established TCP connections. Network and serial connections can be paired by examining the **Target Serial Port(s)** field. The **Status** field describes whether a network connection is established or in the process of being established.

Note: *All counts are from the router's perspective. The **Rx Packets** count reflects packets received from the network, the contents of which are transmitted at the protocol and reflected in the **Tx Chars** field.*

The **Refresh** button will cause the page to be reloaded.

The **Continuous Display** button will cause the browser to continuously reload the page showing the differences in statistics from the last display. **The difference is not a real time rate in bytes or packets per second.**

Protocol Specific Packet Error Statistics

The **Raw Socket** Packet Errors field reflect the number of times that a network message was received and could not be enqueued at the serial port because of output buffering constraints. This is usually symptomatic of a remote peer that uses a higher baud rate or local flow control.

Serial Protocols Trace Menu

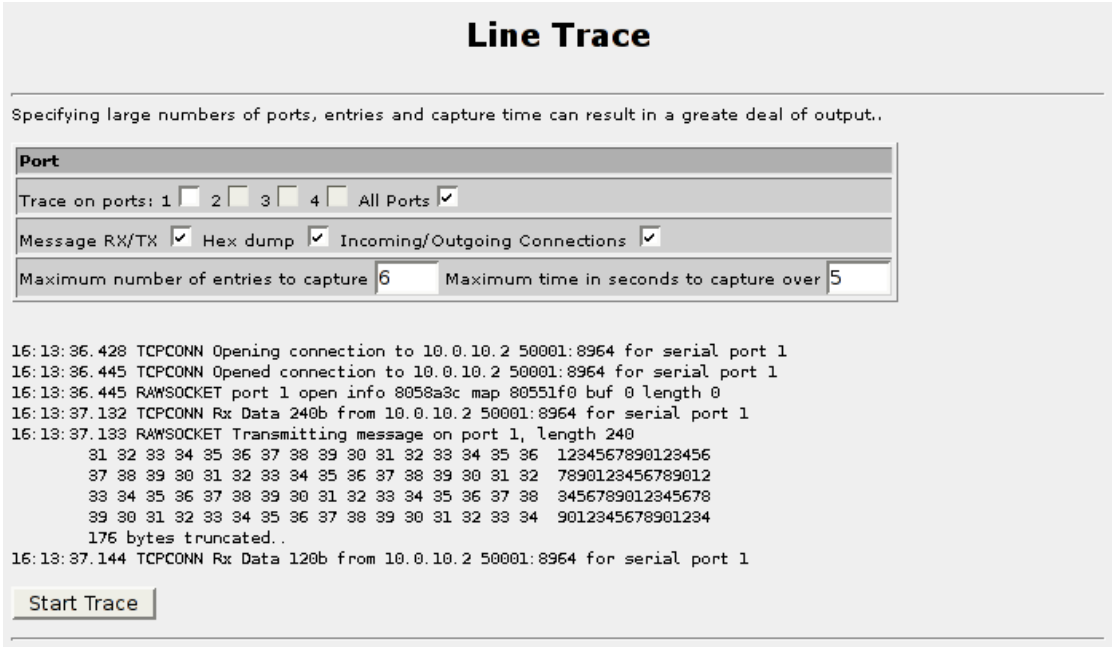

Figure 163: Serial Protocols Trace Menu

This menu displays decoded serial port and network activity.

The desired traffic sources, number of messages and length of time to capture are entered and the **Start Trace** button is pressed. The menu will display up to the provided number of messages waiting up to the specified number of seconds.

The **Trace on ports:** selections feature a list of serial ports with unused entries greyed out. The default is **All Ports**, which selects all ports.

The **Message RX/TX** field allows log entries to be printed for each received or transmitted message, and method of packetization. If the **Hex Dump** field is selected, the first 64 bytes of packet content is displayed.

The **Incoming/Outgoing Connections** field allows regular network level entries such as call connections and received/transmitted messages to be displayed. Note that some unexpected, but unusual, network messages may be displayed if they occur.

Note: Specifying large numbers of ports, entries and capture times can result in a great deal of output. *Specifying a large capture time may require the web page to wait that interval if activity is infrequent.*

Serial Protocols Sertrace Utility

The command line sertrace utility offers the ability to trace the activity of serial ports in real time. A port range may be specified to limit the output to specific ports. The level of traffic to trace and the type of decoding may be specified. The tool may also be used to force the port to transmit an output test message. The following is an example of sertrace use:

```
RuggedRouter:~# sertrace -h

Trace Serial Protocol Server Activity
Usage: sertrace [-dtr] [-p portrange]
       serserver -d protocol decode
       serserver -t tcp level events
       serserver -r raw packet display
       serserver -p ports to capture (e.g 1,3,6-7)
       serserver -s ports to send a test message out on (when 's' + Enter keys are pressed)
In the absence of parameters, all decoding on all ports is provided.

RuggedRouter:~# sertrace -p 1 -s 1

10:56:18.405 TCPCONN Listening on TCP Port 50002 from port 1
10:56:19.944 TCPCONN Connection opened from 10.0.10.236 4991:50002
s
10:56:47.497 RAWSOCKET Transmitting message on port 1, length 44
        74 68 65 20 71 75 69 63 6b 20 62 72 6f 77 6e 20  the quick brown
        66 6f 78 20 6a 75 6d 70 65 64 20 6f 76 65 72 20  fox jumped over
        74 68 65 20 6c 61 7a 79 20 64 6f 67              the lazy dog
10:56:47.545 RAWSOCKET Received message on port 1, length 44 (31ms) by timer
        74 68 65 20 71 75 69 63 6b 20 62 72 6f 77 6e 20  the quick brown
        66 6f 78 20 6a 75 6d 70 65 64 20 6f 76 65 72 20  fox jumped over
        74 68 65 20 6c 61 7a 79 20 64 6f 67              the lazy dog
10:56:47.545 TCPCONN Tx Data from port 1 44b to 10.0.10.236 4991:50002
```

This page intentionally blank

Chapter 20 – Configuring GOOSE Tunnels

Introduction

This chapter familiarizes the user with:

- Configuring GOOSE Tunnels
- Viewing GOOSE Tunnel status and statistics
- Tracing GOOSE activity

IEC61850 GOOSE Fundamentals

IEC61850 is an international standard for substation automation. It is a part of the International Electrotechnical Commission's (IEC) Technical Committee 57 (TC57) architecture for electric power systems.

One feature of IEC61850 is the fast transfer of events. Transfers of Generic Substation Events (*GSE*s) are accomplished through the GOOSE (Generic Object Oriented Substation Event) protocol.

IEC61850 uses Layer 2 multicast frames to distribute its messages and hence, is incapable of operating outside of a switched Ethernet Network. The GOOSE tunnel feature provides a capability to bridge GOOSE frames over a WAN.

GOOSE tunnels provides you with the following features:

- GOOSE traffic is bridged over the WAN via UDP packets.
- One GOOSE traffic source can be mapped to multiple remote router Ethernet interfaces in mesh fashion.
- To reduce bandwidth consumption, GOOSE daemons may be located at each of the "legs" and at the center of a star network. The centrally located daemon will accept GOOSE packets and re-distribute them.
- Statistics reports availability of remote GOOSE daemons, packet counts and Round Trip Time (RTT) for each remote daemon.
- When Virtual Router Redundancy Protocol (VRRP) is employed, GOOSE transport is improved by sending redundant GOOSE packets from each VRRP gateway.

Layer 2 Tunnel Daemon Details

The GOOSE protocol is supported by the Layer 2 Tunnel Daemon. The daemon listens to configured Ethernet interfaces and to the network itself upon a configurable UDP port.

The Media Access Control (MAC) destination address of frames received from Ethernet is inspected in order to determine which GOOSE group they are in. The frames are then encapsulated in network headers and forwarded (with MAC source and destination addresses intact) to the network as GOOSE packets.

IEC61850 recommends that the MAC destination address should be in the range 01:0c:cd:01:00:00 to 01:0c:cd:01:01:ff.

GOOSE Packets received from the network are stripped of their network headers and forwarded to Ethernet ports configured for the same multicast address. The forwarded frames contain the MAC source address or the originating device, and not that of the transmitting interface. The VLAN used will be that programmed locally for the interface and may differ from the original VLAN. The frame will be transmitted with the highest 802.1p priority level (p4).

Packets received from the network will also be forwarded to any other remote daemons included in the group.

Note: Avoid network configurations where the daemons can form a traffic loop. *The simplest such configuration is a triangle network where each daemon forwards to two other routers. Frames arriving at one router will start cycling in clockwise and counterclockwise directions.*

To avoid such "GOOSE storms", frames forwarded to the network are tagged with an initial time to live count. The count is decremented at each relay to the network and prevents the frame from being relayed indefinitely.

Layer 2 Tunnels Main Menu

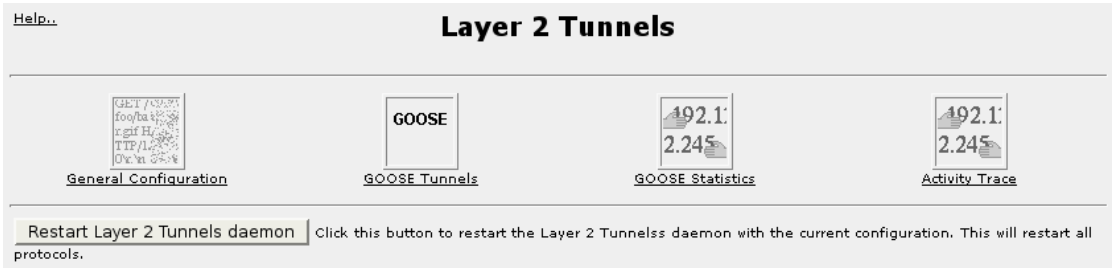

Figure 164: Layer 2 Tunnels Main Menu

Note that the Layer 2 Tunnels daemon is disabled by default and may be enabled via the System folder, Bootup And Shutdown menu.

The **General Configuration** menu changes parameters that apply to all protocols.

The **GOOSE Tunnels** and **GOOSE Statistics** menu configures and display statistics for these tunnels.

The **Activity Trace** menu will provide a protocol trace.

When enabled, any configuration changes may be made to take effect by selecting the **Restart Layer 2 Tunnels daemon** button.

General Configuration Menu

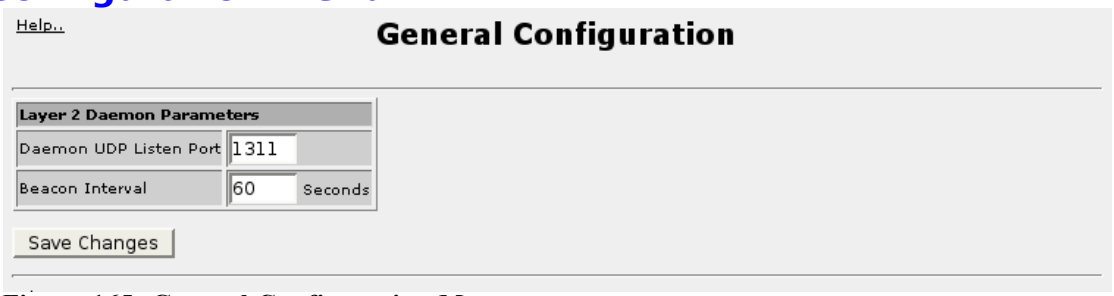

Figure 165: General Configuration Menu

This menu configures the daemon settings.

The **Daemon UDP Listen Port** field configures port used by the daemon to communicate with other daemons.

Note: All Layer 2 Tunnel daemons in the network must use the same port number. *If the router employs a firewall, ensure that a hole is opened for each of the remote daemons using on this port number.*

The Beacon Interval field configures how often a Round Trip Time (RTT) measurement message is sent to each remote peer. The interval takes the values "Off" to disable RTT measurement or a time of 10 – 3600 seconds.

GOOSE Tunnels Menu

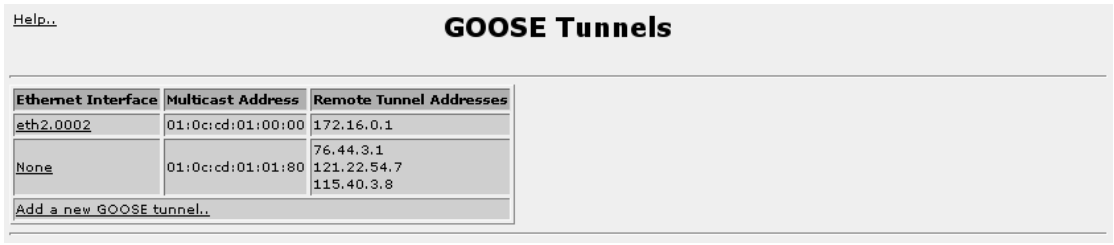

Figure 166: GOOSE Menu

This menu displays GOOSE tunnels. Configure an existing tunnel by following the link under the **Ethernet Interface** field, or add a new tunnel.

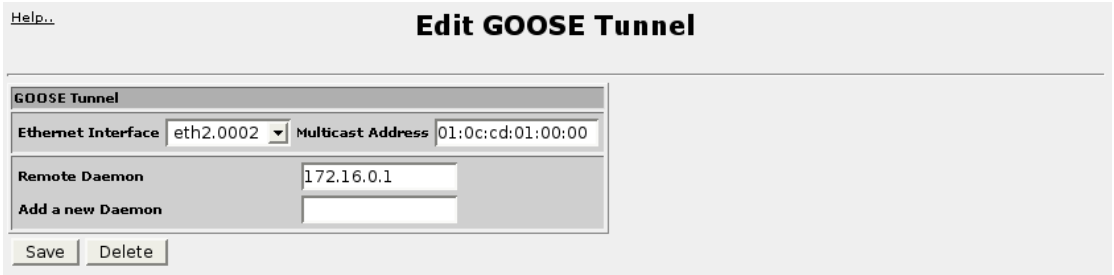

Figure 167: GOOSE Menu

This menu configures a GOOSE tunnel.

The **Ethernet Interface** field configures suitable (i.e. VLAN eligible) interfaces to listen for GOOSE frames upon. You may set this field to "none" if the intent is simply to relay network packets.

The **Multicast Address** field configures the address to listen for.

The **Remote Daemon** and **Add a new Daemon** fields specify the IP addresses of remote daemons.

GOOSE Statistics Menu

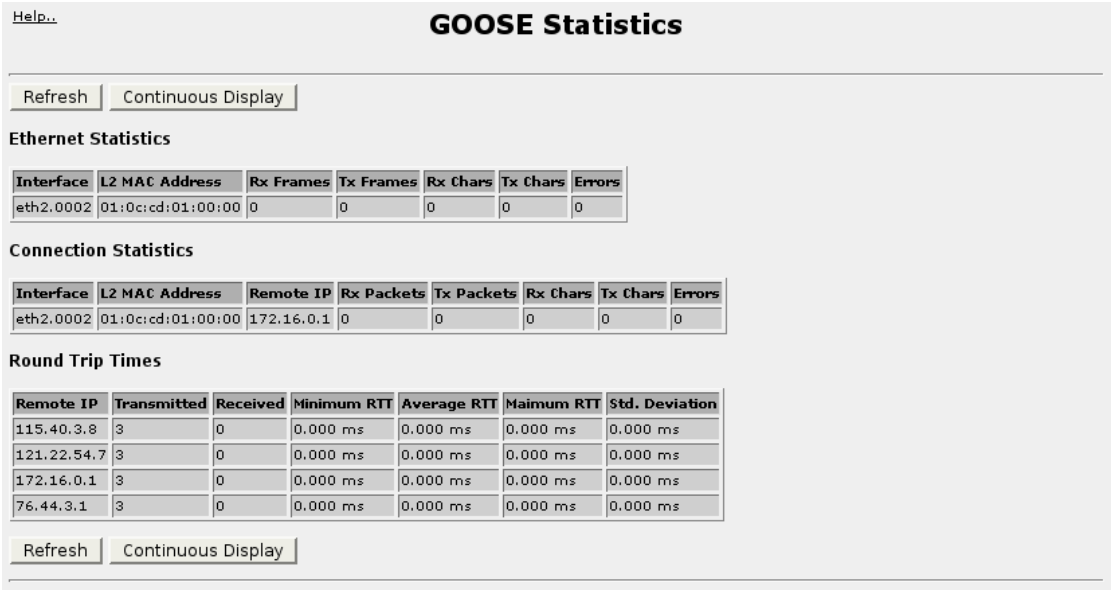

Figure 168: GOOSE Statistics Menu

This menu presents statistics of GOOSE activity at the Ethernet and Network Layers.

The **Ethernet Statistics** table provides a record for each GOOSE tunnel. The number of historical received and transmitted characters as well as errors will be displayed.

The **Connection Statistics** table reflects UDP connections. Network and Ethernet connections can be paired by examining the **L2 MAC Address** field.

Note: *All counts are from the router's perspective. The **Rx Packets** count reflects packets received from the network, the contents of which are transmitted at the protocol and reflected in the **Tx Chars** field.*

The **Round Trip Times** table reflects the measured RTT to each remote daemon. The minimum, average, maximum and standard deviation of times is presented. Wntries with a large difference between the **Transmitted** and **Received** fields indicate potential problems.

The **Refresh** button will cause the page to be reloaded.

The **Continuous Display** button will cause the browser to continuously reload the page showing the differences in statistics from the last display. **The difference is not a real time rate in bytes or packets per second.**

Activity Trace Menu

Figure 169: Activity Trace Menu

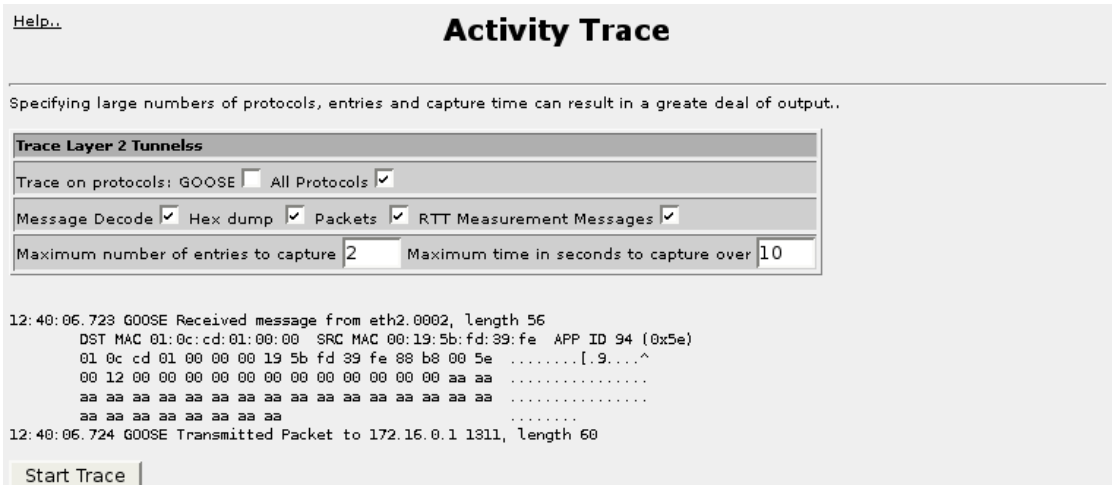

This menu displays decoded GOOSE activity.

The desired traffic sources, number of messages and length of time to capture are entered and the **Start Trace** button is pressed. The menu will display up to the provided number of messages waiting up to the specified number of seconds.

The **Trace on protocols:** selections feature a (all to short) list of protocols with unused entries greyed out. The default is **All Protocols.**

The **Message Decode** field causes received/transmitted frame entries to include protocol specific information. If the **Hex Dump** field is selected, the first 64 bytes of packet content is displayed.

The **Packets** field causes received/transmitted packet entries to be displayed.

The **RTT Measurement** field displays Beacon messages used for RTT measurement.

Note: Specifying large numbers of ports, entries and capture times can result in a great deal of output. *Specifying a large capture time may require the web page to wait that interval if activity is infrequent.*

This page intentionally blank

Chapter 21 - Configuring The DHCP server

Introduction

This chapter familiarizes the user with:

- DHCP Server Configuration
- Use of Option 82

DHCP Fundamentals

Dynamic Host Configuration Protocol (DHCP) is a method for centrally and consistently managing IP addresses and settings for clients, offering a variety of assignment methods. IP addresses can be assigned based on the Ethernet MAC address of a client, sequentially, or by using port identification provided by a DHCP relay agent device.

DHCP Network Organizations

The information to assign addresses in DHCP is organized to deal with clients at the host, group, subnet, pool and shared network level.

Hosts entries assign specific settings to a client based on its Ethernet MAC address.

Groups allow identical settings to be created for a group of hosts, making it simpler to manage changes to the settings for all the hosts contained within the group. *Groups contain hosts*.

Pools contain ranges of IP addresses to hand out to clients with access rules to determine which clients should receive addresses from that pool.

Subnets control settings for each subnet that DHCP serves. A subnet can include a range of IP address to hand out to clients. Only one subnet can contain dynamic IP address ranges without any access restrictions on any given physical port since DHCP doesn't know which subnet a client should belong to when the request is received. *Subnets contain groups, pools and hosts.*

Shared networks are used when multiple subnets should be served by a single physical port. This applies both when using a DHCP relay agent connected to the port with additional subnets behind the relay agent, or when multiple virtual networks exist on one physical interface. Each subnet then gets its own subnet definition inside the shared network rather than at the top level. *Shared networks contain subnets, groups and hosts.*

DHCP Client Options

The following options apply to single hosts, subnets of hosts, pools (potentially discontinuous ranges of addresses), shared networks (a single physical networks for which distinct subnets of hosts coexist and request addresses) and groups. The meaning of each option is the same in each case, while the type of target determines which clients it applies to.

In DHCP settings at a more specific level overrides higher levels. For example you can configure a DNS server for all clients, the create a group that overrides the setting. This allows defaults to be set at a high level to apply to most clients, while exceptions can be places just where they are needed. Many settings are only supported by certain specific types of clients, and are ignored by the majority of clients.

Basic options you should pay attention to include:

- Address ranges: The range of addresses to use for dynamic IP clients.
- Default lease time: The default length of leases assigned to clients, if the client doesn't request a lease length.
- Maximum lease time: The maximum length of leases allowed to clients. If a client requests a higher value it will be refused.
- Client hostname: The hostname the client should use.
- Default routers: The default gateway the client should use.
- Domain name: The DNS domain name the client should use.
- DNS servers: The IPs of the DNS servers the client should use.
- NTP servers: The IPs of the NTP servers the client should use.
- Static routes: Static routes the client should use.
- Time servers: The IPs of the time servers the client should use.

Lesser used client options include:

- Subnet mask: The subnet mask the client should use. Rarely needed.
- Broadcast address: The broadcast address the client should use. Rarely needed.
- Log servers: The IPs of the LOG servers the client should use.
- Swap server: The IP of the swap server the client should use. Normally only used for diskless network booted clients.
- Root disk path: The path the client should use for its root device. Normally only used for diskless network booted clients.
- NIS domain: The NIS domain the client should use.
- NIS servers: The IPs of the NIS server the client should use.
- Font servers: The IPs of the font servers the client should use. Normally only used for X terminals.
- XDM servers: The IPs of the XDM servers the client should use. Normally only used for X terminals.
- NetBIOS name servers: The IPs of the Netbios name servers the client should use.
- NetBIOS node type: The NetBIOS name resolution method the client should use.
- NetBIOS scope: The NetBIOS scope the client should use.
- Time offset: The offset from a time server the client should be using.
- Custom options allows you to add additional DHCP options required by a client.

BOOTP and Dynamic DNS related options include:

- Boot filename: The filename the client should request from a tftp server to boot from. This only applies to network booted clients.
- Boot file server: The IP address of the tftp server to boot from. This only applies to network booted clients.
- Server name: The hostname of the boot server. This only applies to network booted clients.
- Lease length for BOOTP clients: How long the IP assigned to a BOOTP client should be considered valid.
- Lease end for BOOTP clients: Cut off date for all BOOTP client leases.
- Dynamic DNS enabled: Should DNS information be updated on the DNS server when a client receives an IP address.
- Dynamic DNS domain name: The domain name to update dynamic DNS information in.
- Dynamic DNS hostname: Use the specified hostname for clients, or use the hostname supplied by the client.
- Dynamic DNS reverse domain: The reverser DNS domain to update dynamic information in for the reverse DNS entry.
- Dynamic DNS reverse domain: The reverser DNS domain to update dynamic information in for the reverse DNS entry.

Lesser used DHCP server configurations include

- Allow unknown clients: Should DHCP accept requests from clients it has never seen before or only from clients that have already received leases in the past.
- Server is authoritative: If the server is authoritative, it will send deny messages to any client which tries to renew a lease which the server knows the client shouldn't have.
- Option 82 Support.

Option 82 Support with Disable NAK

If DHCP relay clients (option 82 clients) are used on the same subnet as the DHCP server, some clients will immediately try to renew a lease right after receiving it by requesting a renewal directly from the DHCP server. Since the DHCP server is only configured to provide that lease through a relay agent with the right option 82 fields added, the server will send the client a NAK to disallow use of the lease. Enabling this option disables this reject message, so that the renewal request that the DHCP relay agent sends a moment later (which the DHCP server accepts since it has the right option 82 fields added) will be the only message for which the client receives a reply. If the DHCP server and clients are not on the same subnet, this option is not required. The meaning of the value of many fields depends on the client's interpretation of the field, so the actual meaning of a field is determined by the client. See the documentation of the client to determine what values are required by the client for special options.

Example DHCP Scenarios And Configurations

Single Network With Dynamic IP Assignment

In this example the eth1 interface is provided with IP address 192.168.1.1/24 while addresses 192.168.1.101 through 192.168.1.200 are assigned to the clients. The router serves as the default gateway.

1) Enable eth1 in the 'Edit Network Interfaces' menu.
2) Click 'add a subnet', and configure it for network address 192.168.1.0 with netmask 255.255.255.0.
3) Set the assigned address range to 192.168.1.101 - 192.168.1.200.
4) Click 'Create' then edit the subnet just created and click 'Edit Client Options'.
5) Set default routers to 192.168.1.1 and save.
6) Restart the DHCP server or apply changes.

Single Network With Static IP Assignment

In this example the eth1 interface is provided with IP address 192.168.1.1/24.

Assign address 192.168.1.101 to a DHCP client with MAC 00:11:22:33:44:01.

Assign address 192.168.1.102 to a DHCP client with MAC 00:11:22:33:44:02.

Assign address 192.168.1.103 to a DHCP client with MAC 00:11:22:33:44:03.

The router serves as the default gateway.

1) Enable eth1 in the 'Edit Network Interfaces' menu.
2) Click 'add a subnet', and configure it for network address 192.168.1.0 with netmask 255.255.255.0.
3) Click 'Create' then edit the subnet just created and click 'Edit Client Options'.
4) Set default routers to 192.168.1.1 and save it.
5) Click 'add a new host'.
6) Set the hardware address to Ethernet 00:11:22:33:44:01 and the fixed IP to. 192.168.1.101. Assign the client a hostname as well.
7) Click 'Create'.
8) Repeat steps 5) through 7) for the other hosts with the appropriate address, MAC and hostname for each client.
9) Restart the DHCP server or apply changes.

Single Network With Option82 Clients On One Switch

In this example the eth1 interface is provided with IP address 192.168.1.1/24

A a switch connected to eth1 and uses address 192.168.1.2/24.

The switch port 1 is connected to the router while its ports 2 through 8 provide DHCP relay support. The switch has its DHCP relay server address set to router's address 192.168.1.1. The switch has all ports in VLAN 1. The switch base MAC address is 00:0A:DC:11:22:00.

Assign a client at switch port 2 address 192.168.1.102.

Assign a client at switch port 3 address 192.168.1.103.

Assign multiple clients at switch port 4 dynamic addresses 192.168.1.151 through 192.168.1.200.

The router serves as the default gateway.

1) Enable eth1 in the 'Edit Network Interfaces' menu.
2) Add a new subnet, and configure it for network address 192.168.1.0 with netmask 255.255.255.0.
3) Enable the 'Disable NAK of option82 clients for this subnet?' option to prevent confusing some DHCP clients due to the client being on the same network as the DHCP server and the DHCP relay agent (the switch).
4) Save it then edit the subnet just created and click 'Edit Client Options'.
5) Set default routers to 192.168.1.1 and save it.
6) Click 'add an address pool' to the subnet.
7) Set the address range to 192.168.1.102 to 192.168.1.102.
8) Click 'Create'.
9) Edit the pool by clicking on the link for the pool with address range 192.168.1.102 - 192.168.1.102.
10) Click 'add an option82 client'.
11) Give the client a unique alpha numeric name (for example client0102).
12) Set the remote id to the switch MAC address (00:0A:DC:11:22:00 in this case).
13) Set the circuit id to the switches circuit id identifier to the port (00:01:00:02 for VLAN 1 port 2 on a RuggedCom switch).
14) Click 'Create'.
15) Click 'Save'.
16) Repeat steps 6) through 15) for clients 192.168.1.103 changing the pool address range and circuit id.
17) Repeat steps 6) through 15) for port 4 using the address range 192.168.1.151 to 192.168.1.200 and the circuit id for port 4.
18) Restart the DHCP server or apply changes.

Multiple Subnets On Separate VLANs Using Option82 On One Switch

In this example the eth1 interface is provided with IP address 192.168.1.1/24

A switch connected to eth1 and using address 192.168.1.2/24

The switch port 1 is connected to the router while its ports 2 through 8 provide DHCP relay support. The switch has its DHCP relay server address set to router's address 192.168.1.1. The switch has all ports in VLAN 1. The switch base MAC address is 00:0A:DC:11:22:00.

The switch port 2 is on vlan2 using subnet 192.168.2.0/24 and should assign addresses 192.168.2.101 to 192.168.2.200 and default gateway 192.168.2.1.

The switch port 3 is on vlan3 using subnet 192.168.3.0/24 and should assign addresses 192.168.3.101 to 192.168.3.200 and default gateway 192.168.3.1.

The switch port 4 is on vlan4 using subnet 192.168.4.0/24 and should assign addresses 192.168.4.101 to 192.168.4.200 and default gateway 192.168.4.1.

1) Enable eth1 in the 'Edit Network Interfaces' menu.
2) Add a new subnet, and configure it for network address 192.168.1.0 with netmask 255.255.255.0.
3) Save it.
4) Add a new shared network.
5) Name the shared network (for example "eth1") and select the subnet 192.168.1.0 to be included in the shared network.
6) Save it.
7) Edit the shared network again.
8) Add a new subnet, and configure it for network address 192.168.2.0 with netmask 255.255.255.0
9) Save the new subnet and then save the shared network settings.
10) Edit the subnet just created and click 'Edit Client Options'.
11) Set default routers to 192.168.2.1 and save it.
12) Click 'add an address pool' to the subnet.
13) Set the address range to 192.168.2.101 to 192.168.2.200.
14) Click 'Create'.
15) Edit the pool by clicking on the link for the pool with address range 192.168.2.101 - 192.168.2.200.
16) Click 'add an option82 client'.
17) Give the client a unique alpha numeric name (for example subnet0102).
18) Set the remote id to the switch MAC address (00:0A:DC:11:22:00 in this case).
19) Set the circuit id to the switches circuit id identifier to the port (00:02:00:02 for VLAN 2 port 2 on a RuggedCom switch).
20) Click 'Create'.
21) Click 'Save'.
22) Repeat steps 8) through 20) for vlan3 through vlan4 changing the subnet, default routers, pool address range and circuit id for each vlan.
23) Restart the DHCP server or apply changes.

DHCP Server Main Menu

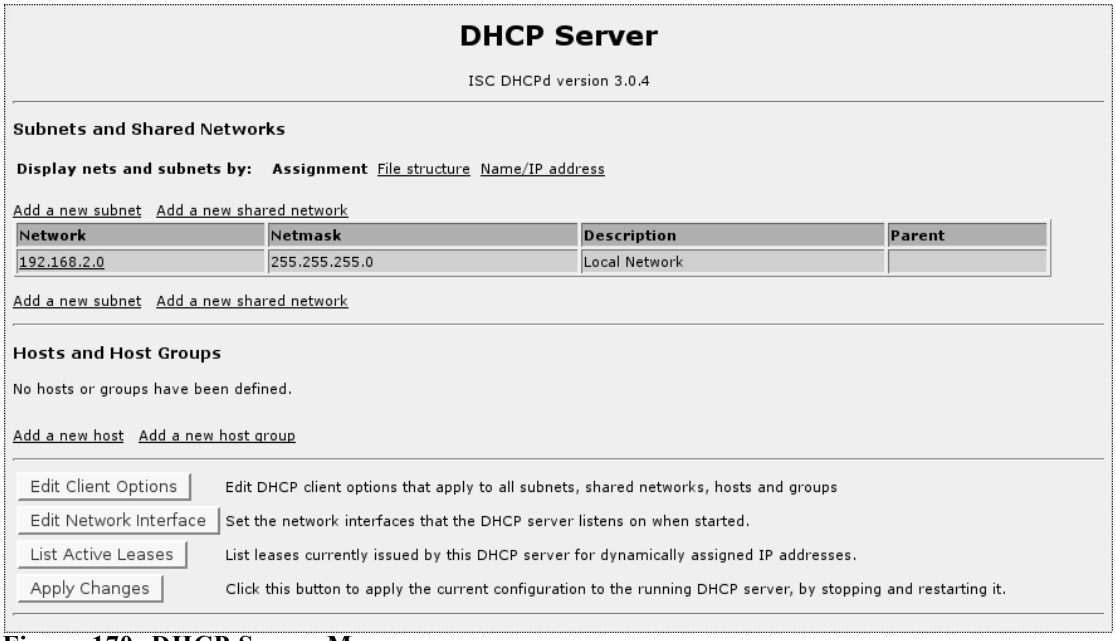

Figure 170: DHCP Server Menu

The DHCP Server main menu shows the subnets configured for DHCP, as well as any groups and hosts. New subnets, groups and hosts can be added, and existing entries can be edited (and optionally deleted).

The **Edit Client Options** button allows you to set global client settings for the DHCP server. Settings made here apply to all clients unless overridden at a lower level in the configuration.

The **Edit Network Interface** button allows you to select which interfaces DHCP should listen for DHCP requests on. Note that you must also have a subnet matching the IP address of the selected interface configured in DHCP in order to actually have DHCP listen for requests on a port.

The **List Active Leases** button shows you which dynamic IP leases are currently assigned to clients.

The **Start Server** button starts the server to check the configuration. To permanently enable DHCP you should enable it in the bootup and shutdown menu.

The **Apply Changes** button applys new settings to the running DHCP server. Use this after making any changes to the configuration.

DHCP Shared Network Configuration

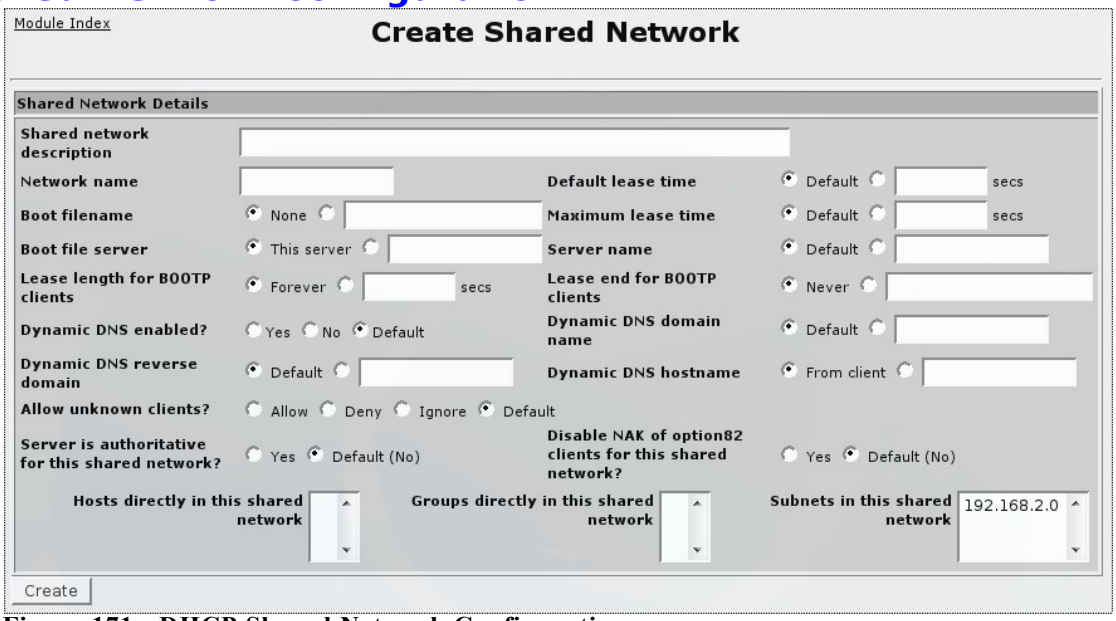

Figure 171: DHCP Shared Network Configuration

The settings specific to the Shared network menu are the Shared network description and Network name.

The **Shared network description** field is used to describe the shared network as desired.

The **Network name** field is a unique name to assign to the shared network. It could be the name of the interface the shared network is on, for example.

Within a shared network you can great subnets, hosts, and groups of hosts.

DHCP Subnet Configuration

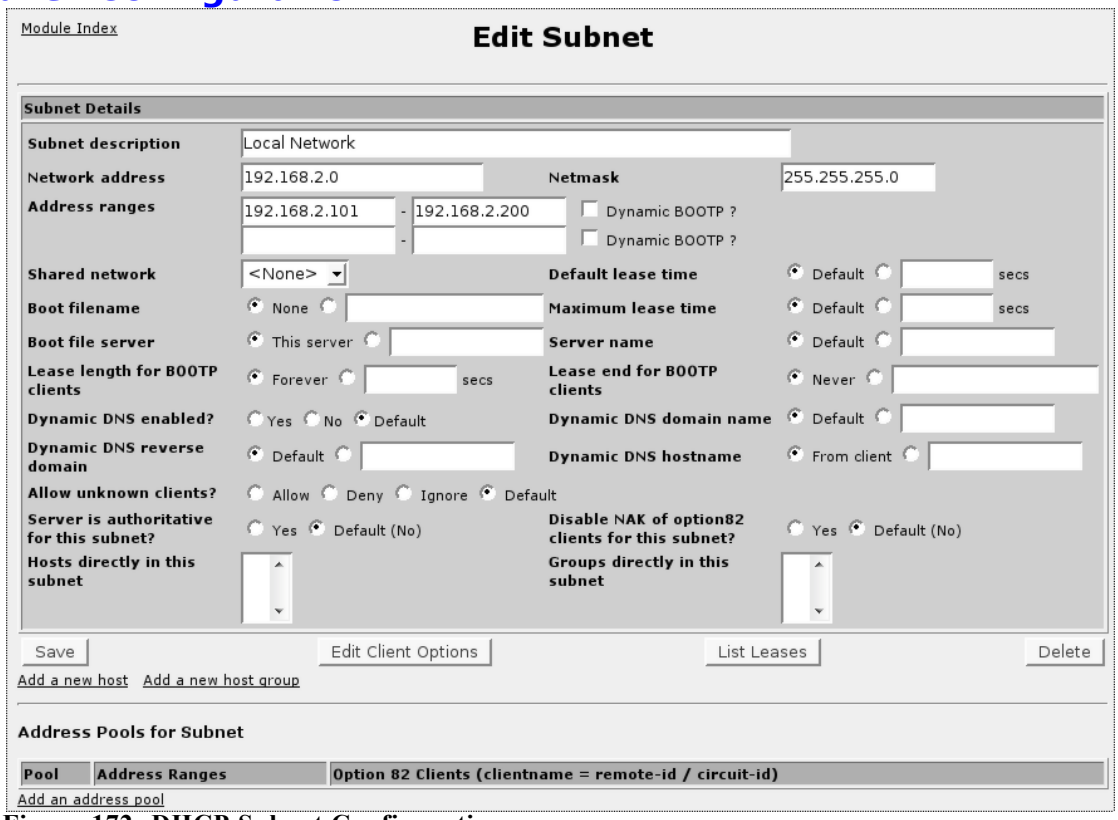

Figure 172: DHCP Subnet Configuration

The settings specific to the Subnet menu are the subnet description, Network address and mask.

The **Subnet description** field is used to describe the subnet as desired.

The **Network address and Netmask fields** of the subnet help to specify the span of assigned addresses.

Within a subnet you can great hosts, groups of hosts, and address pools.

DHCP Group Configuration

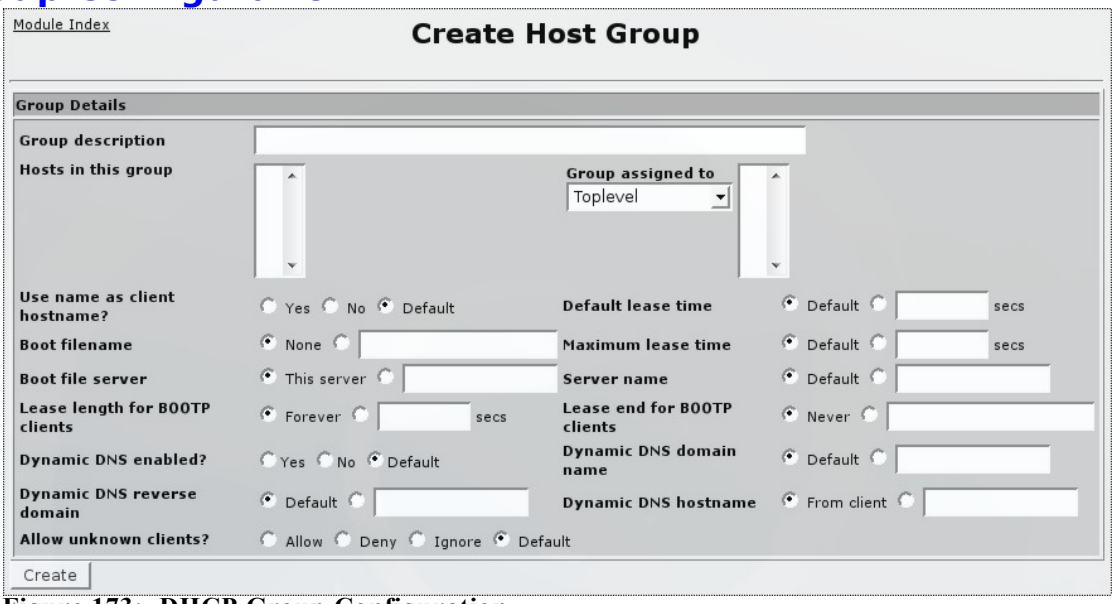

Figure 173: DHCP Group Configuration

The settings specific to the Group menu are the group description and Use name as client hostname fields.

The **Group description** field is used to describe the group as desired.

The **Use name as client hostname** field determines whether host entries should use the hosts entry name as the client hostname to provide to the client.

Within a group you can create hosts.

DHCP Host Configuration

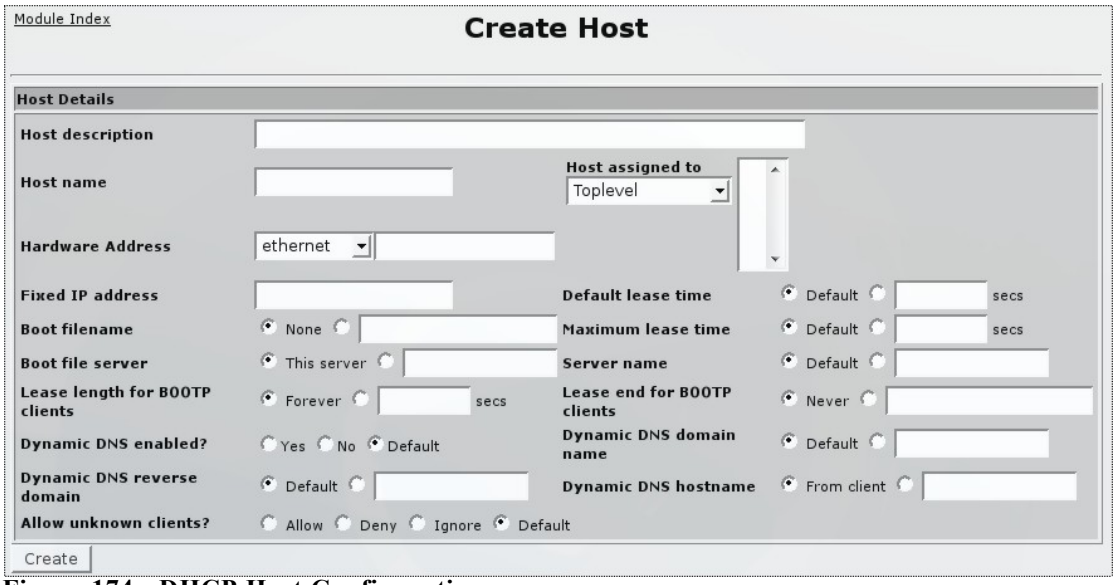

Figure 174: DHCP Host Configuration

The **Host description** field is used to describe the host as desired.

The **Host name** field is the unique name to refer to the host within the DHCP configuration.

The **Hardware address** field is the Ethernet MAC of the client associated with the host entry.

The **Fixed IP address** field is the IP to assign to the matching client.

DHCP Pool Configuration

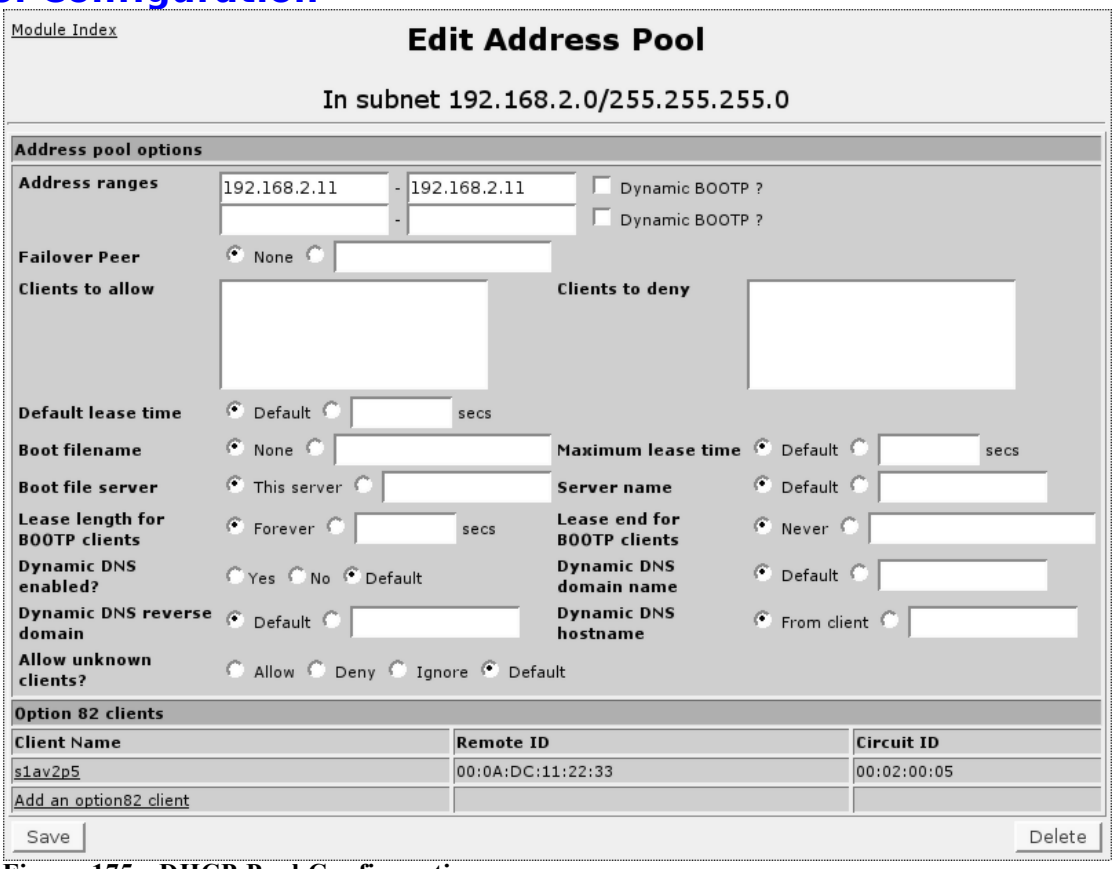

Figure 175: DHCP Pool Configuration

The settings specific to the Address Pool menu are the Failover peer and Clients to allow/deny.

The **Failover peer** field is the IP address of a DHCP peer server if a fail over pool is created.

The **Clients to allow/deny** field can be used to control which clients can get IP address from the pool. See documentation for dhcpd3 for syntax and allowed values. Very rarely needed. The Allow unknown clients setting already handles the most common use of this option.

Chapter 22 – Configuring NTP

Introduction

This chapter familiarizes the user with:

- Enabling/Disabling NTP
- Setting servers and peers
- Setting generic NTP options
- NTP Tools

NTP Fundamentals

NTP (Network Time Protocol) is an Internet protocol used to synchronize the clocks of computers to some time reference. Variants of NTP such as SNTP (Simple NTP, a reduced functionality NTP) and XNTP (Experimental NTP) exist. NTP itself is available in versions 3 and 4 (the RuggedRouter includes version 4).

NTP is a fault-tolerant protocol that allows an NTP daemon program to automatically select the best of several available time sources, or reference clocks, to synchronize to. Multiple candidates can be combined to minimize the accumulated error. Temporarily or permanently wrong time sources are detected and avoided.

The NTP daemon achieves synchronization by making small and frequent changes to the router hardware clock.

The NTP daemon operates in a **client-server mode**, both synchronizing from **servers** and providing synchronization to **peers**.

If NTP has a number of servers to choose from, it will synchronize with the lowest stratum server. The stratum is a measure of the number of servers to the (most highly accurate) reference clock. A reference clock itself appears at stratum 0. A server synchronized to a stratum *n* server will be running at stratum *n* + 1.

You will generally configure lower stratum NTP hosts as servers and other NTP hosts at the same stratum as peers. If all your configured servers fail, a configured peer will help in providing the NTP time. It is generally a good idea to configure one at least one server and peer.

The NTP daemon will know about the NTP servers and peers to use in three ways.

- It can be configured manually with a list of servers to poll from,
- It can be configured manually with a list of peers to send to,
- It can look at advertisements issued by other servers on multicast or broadcast addresses.

Note that if multicasting or broadcasting is used, it is strongly recommended to enable authentication unless you trust all hosts on the network.

NTP uses UDP/IP packets for data transfer because of the fast connection setup and response times UDP offers. The NTP protocol uses port UDP port 123. Note that if your router employs a firewall and acts as a client it must open UDP port 123. Additionally, if the router acts as a server the firewall must allow connection requests on port 123 as well.

The NTP Sanity Limit

NTP changes the system through “stepping” and “drifting”. Stepping is a sudden change of time whereas drifting is a slow gradual time change.

NTP will step the system time when its starts. This is almost always at boot time. Stepping the time afterwards can cause protocols (such as OSPF) that rely upon accurate real time to fail. The router deals with this problem by restarting these protocols if they are running when NTP restarts.

After booting, NTP uses drifting to achieve synchronization by making small and frequent changes to router hardware clock. If the synchronizing server's clock differs from the hardware clock by more than 1000 seconds, the NTP daemon construes a major problem and terminates.

Usually, NTP will succeed in synchronizing the clock at boot time. If it fails to synchronize the clock (perhaps due to a downed WAN link), the NTP daemon may terminate. The router, however, will note the termination and will restart the NTP daemon.

NTP And The Precision Time Protocol Card

If the router is equipped with a Precision Time Protocol card, NTP will treat the Global Positioning System signals received from the card (when GPS locks) as a stratum 0 reference clock. The router will always preferentially use this reference above all others.

Included With NTP

Your RuggedRouter software includes the ntpq, ntpdc, ntptrace and ntp-keygen command line utilities. The ntpq utility program can be used to monitor the NTP daemon operations and determine how well it is running. The ntpdc utility program is used to query the NTP daemon about its current state and to request changes in that state. The ntptrace utility is a utility trace a chain of NTP servers back to the primary source.

The ntp-keygen utility can be used to generate secure public keys for authentication.

NTP Server Main Menu

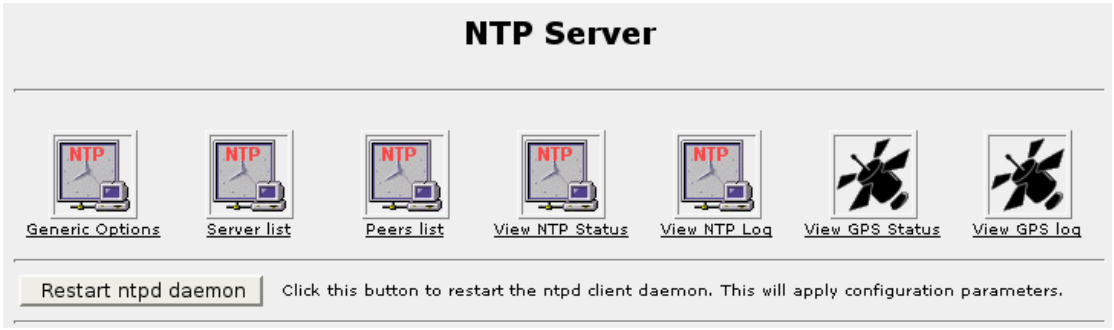

Figure 176: NTP Server

Note that the NTP server is disabled by default and may be enabled via the System folder, Bootup And Shutdown menu. When enabled, any configuration changes may be made to take effect by selecting the **Restart ntpd daemon** button. The **View GPS Status** and **View GPS log** sub-menus appear if the router is equipped with a Precision Time Protocol card.

Generic Options

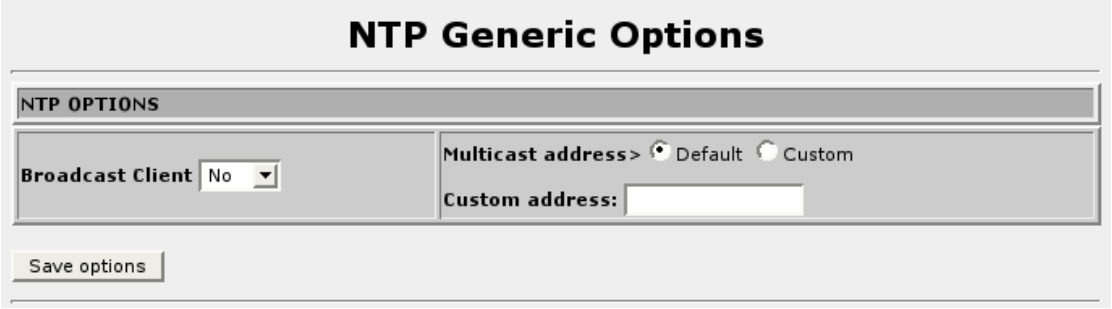

Figure 177: NTP Generic Options

Set the **Broadcast Client** option to "Yes" if you wish to act on NTP broadcast messages.

The default multicast address used for NTP is 224.0.1.1. Select a custom multicast address with the **Custom address** field if you wish to use a different addresses.

Servers Configuration

Figure 178: NTP Server List

The servers under the **IP address** column are used as primary synchronization devices. Clicking on a link will allow you to edit that server.

By default the router includes the links pool.ntp.org and 127.127.1.0. The pool.ntp.org address selects a random low stratum server from a pool of ntp servers on the Internet. The 127.127.1.0 is known as a pseudo-address and points to the local hardware clock of the router. This address allows the router to act as a high stratum NTP server to locally connected devices while in a standalone mode.

If you are operating in a private network you will want to delete both of these addresses and substitute that of a locally known low stratum server.

The **Version** field indicates the version of the NTP protocol used to communicate with this host. Change this only if it is known that the host requires a version other than 4.

The **Key** field provides an authentication key associated with this host.

The **Preferred** field determines whether this host is preferred over other hosts in the list.

The **Check** field link leads to a page that displays the result of an NTP query to this host. Use this feature to determine if the configured host is active.

Peers Configuration

This menu allows you to enter and edit peers. Peers are NTP servers of the same stratum as the router, and are useful when contact is lost with the hosts in the NTP servers menu.

The per-peer configuration information is as described in the previous menu.

Viewing The NTP Status

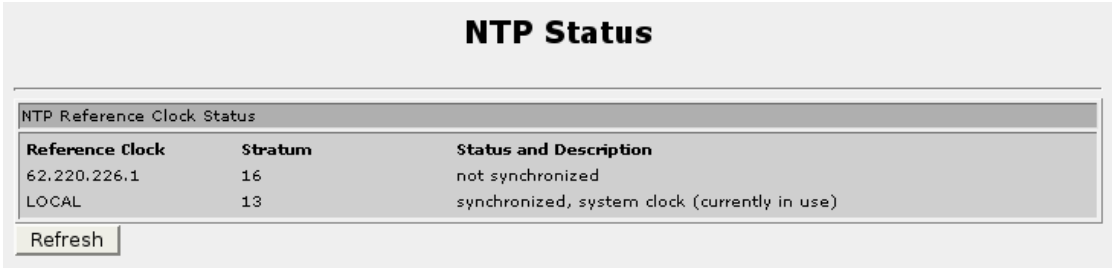

Figure 179: NTP Status

The NTP Status menu displays possible sources and currently used reference clocks

Viewing The NTP Log

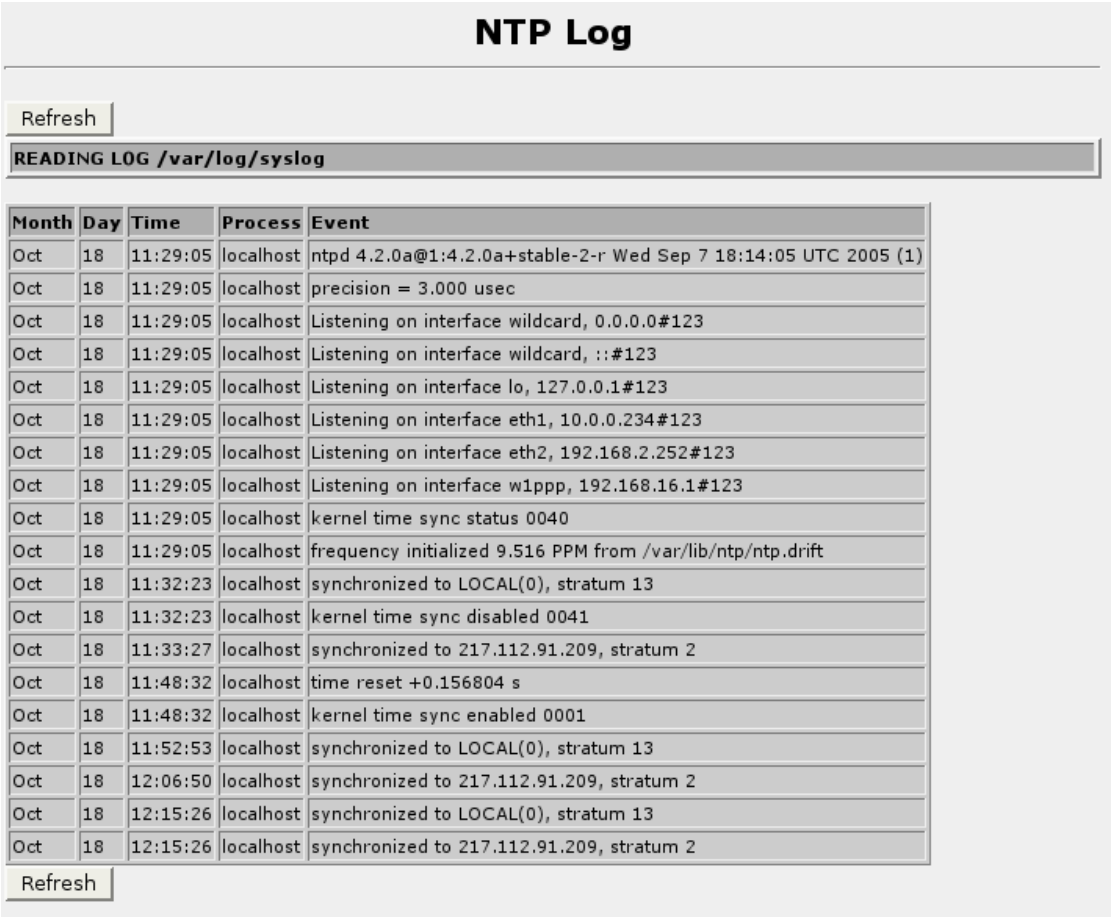

Figure 180: NTP Log

The NTP Log menu displays the log of recent NTP events.

Viewing The GPS Status

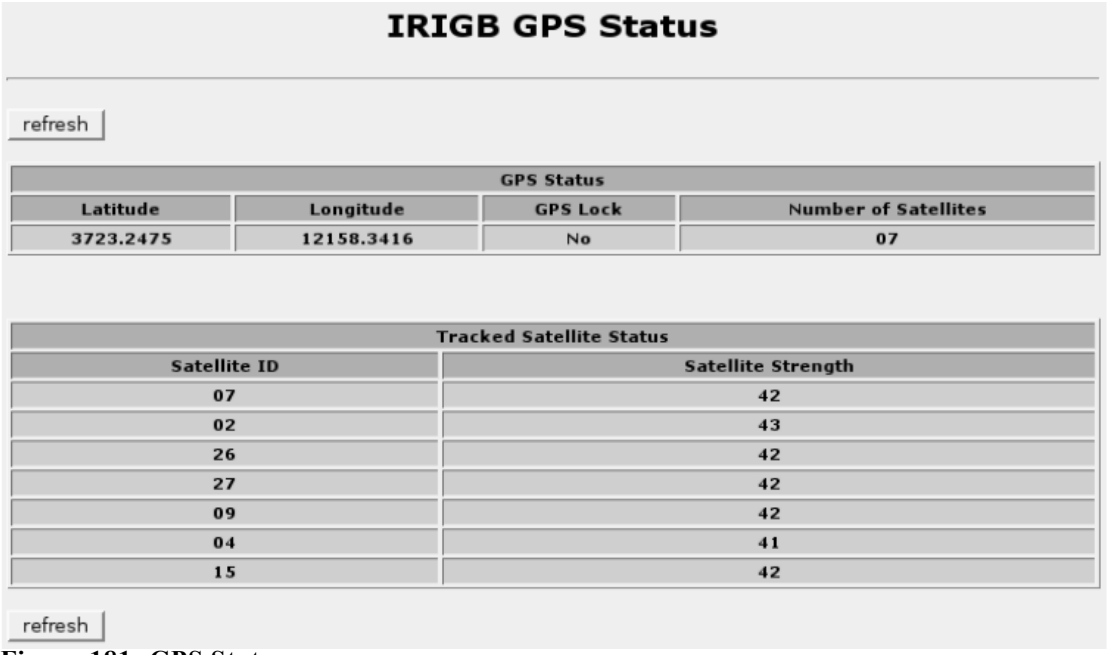

Figure 181: GPS Status

If the router is equipped with a Precision Time Protocol card, this page will shows the status of the GPS module.

The **Latitude** and **Longitude** fields show the current position of the GPS antenna.

The **GPS Lock** field show the GPS lock status.

The **Number of Satellites** shows how many satellites are currently being tracked by the GPS module.

The **Tracked Satellite Status** table shows the ID and signal strength of tracked satellites.

Viewing The GPS Log

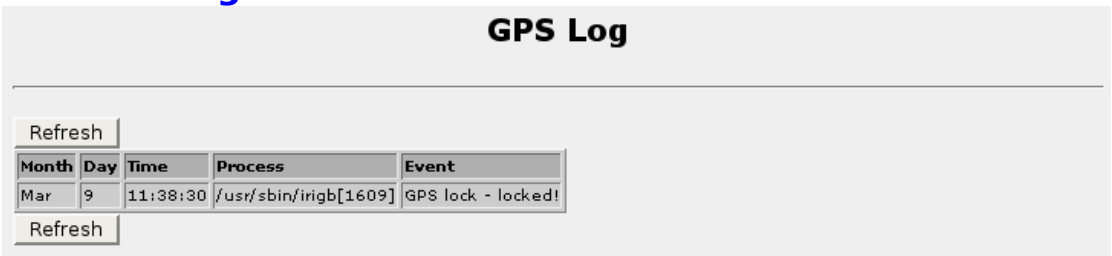

Figure 182: GPS Log

The GPS Log menu displays the log of recent GPS events.

Chapter 23 – Configuring SSH

Introduction

This chapter familiarizes the user with:

- Configuring SSH Authentication
- SSH Networking And Access Control
- Setting SSH Client Options

SSH Fundamentals

Secure Shell is a program to allow logging into another host, to remotely execute commands, and to move files from one machine to another. It provides strong authentication and secure communications over insecure channels. The program that accepts the SSH client's connection is an SSH server. The SSH server can be programmed to enforce conditions to increase security. These conditions can be imposed upon specific hosts or upon all hosts in general.

SSH has had two major revisions of the protocol upon which it is based, SSH v1 and v2. SSH v1 relied upon the RSA authentication scheme, while SSH v2 relies upon RSA or DSA. SSH v1 is known to be insecure and should not be used.

SSH operates upon TCP port 23 by default. Open this port if you use a firewall.

SSH also provides TCP forwarding, a means to forward otherwise insecure TCP traffic through SSH Secure Shell.

Included With SSH

Your RuggedRouter software includes scp, an SSH utility to perform secure copying of files and directories over the network.

If you decide to create you own user accounts, the ssh-keygen utility can be used to populate the account with SSH keys.

SSH Main Menu

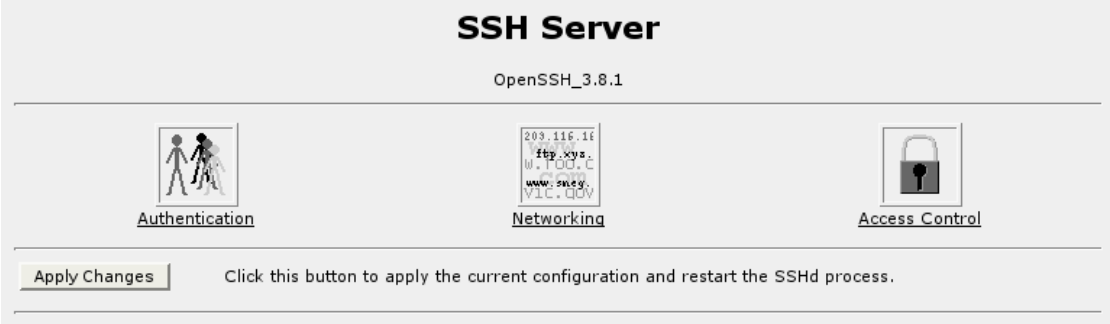

Figure 183: SSH Server

Note that the SSH server is enabled by default and may be disabled via the System folder, Bootup And Shutdown menu. When enabled, any configuration changes may be made to take effect by selecting the **Apply Changes** button.

Authentication

Figure 184: SSH Server Authentication Menu

The **Allow authentication by password** field determines whether to allow clear text tunneled passwords. If set to Yes, the user will be allowed to enter a password for authentication if validation cannot be done using a public key.

The **Permit logins with empty passwords** field (when password authentication is allowed) specifies whether the server allows login to accounts with empty passwords.

The **Allow RSA authentication field s**pecifies whether pure RSA authentication is allowed. If this is set to "No", users will always need to enter their password even if their public key has been set up.

Networking

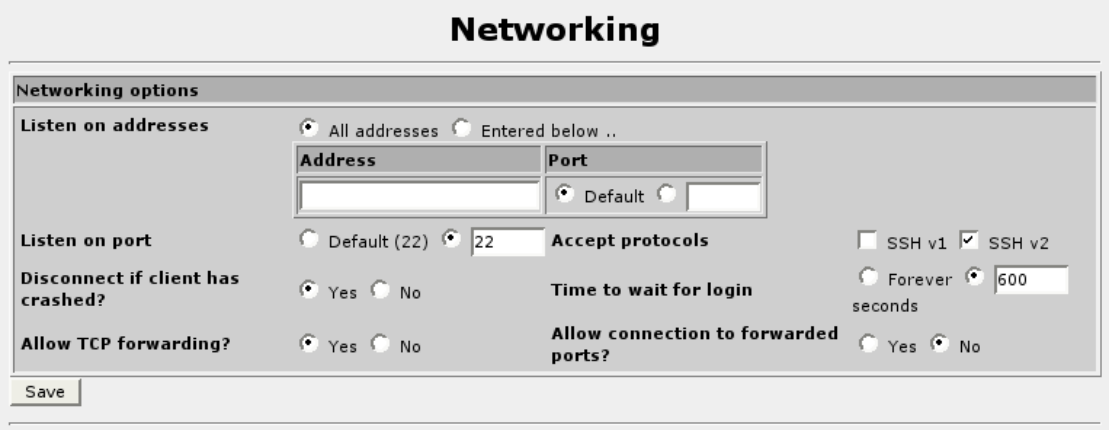

Figure 185: SSH Server Networking

The **Listen on addresses** fields determine an IP addresses and port upon which SSH will accept a connection.

The **Listen on port** field determines the port number SSH will listen on, assuming **Listen on addresses** is set to "All addresses".

The **Accept Protocols** fields determine which versions of SSH will be allowed.

The **Disconnect if client has crashed** field determines whether the SSH server should periodically check to see if the client is still alive.

The **Time to wait for login** field determines the maximum time from a connection request until login completes, after which the client will be disconnected.

The **Allow TCP forwarding** field specifies whether TCP forwarding is permitted. If this option is set, clients on a remote network can tunnel TCP connections to machines on the RuggedRouter's network.

The **Allow connection to forwarded ports** field specifies whether remote hosts on the client network are allowed to connect to ports forwarded for the client.

Access Control

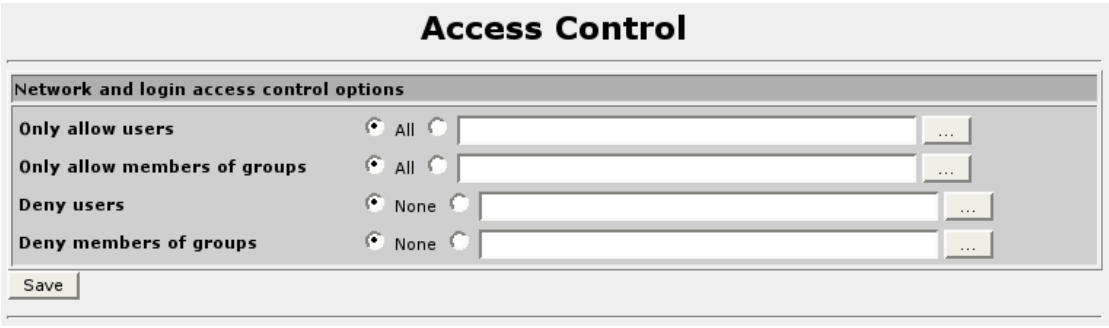

Figure 186: SSH Server Access Control

The **Only allow users** field specifies the users allowed to connect by SSH. The specification can be a list of user name patterns, separated by spaces. Login is allowed only for user names that match one of the patterns. '*' and '?' can be used as wild cards in the patterns. Only user names are valid, a numerical user ID is not recognized. By default, login is allowed for all users. If the pattern takes the form USER@HOST then USER and HOST are separately checked, restricting logins to particular users from particular hosts.

The account selector () button can be user to build up a list of allowable users.

The **Only allow members of groups** field specifies the "group" (in the Unix sense) of users allowed to connect by SSH. The specification can be followed by a list of group name patterns, separated by spaces. If specified, login is allowed only for users whose primary group or supplementary group list matches one of the patterns. '*' and '?' can be used as wild cards in the patterns. Only group names are valid, a numerical group ID is not recognized. By default, login is allowed for all groups.

The account selector () button can be user to build up a list of allowable groups.

The **Deny users** and **Deny members of groups** fields specify users and groups to deny connections to.

Chapter 24 – Configuring IRIGB And IEEE1588

Introduction

This chapter familiarizes the user with:

- IEEE 1588 Configuration
- IRIGB Configuration
- Viewing IRIGB and IEEE1588 Status

IEEE1588 Fundamentals

The IEEE 1588 working group Precise Timing Protocol (PTP) standard details a method of synchronizing a clocks over networks, including Ethernet. The RuggedRouter provides a special hardware assisted PTP capability as provided by the RuggedCom PTP card. When used in conjunction with the cards Global Positioning System (GPS) receiver, the router can provide nanosecond accuracy via IEEE1588.

Additionally, IEEE1588 may be used (in GPS failure situations) to synchronize to a remote source and provide accurate time for IRIG-B.

PTP Network Roles

The IEEE 1588 standard describes regular *clocks* as devices having a single PTP port that can issue and receive PTP messages. PTP *boundary clocks* are clocks have have multiple PTP ports, offering the ability to serve time to more than one subnet at a time. The RuggedRouter can serve as a regular clock and communicate with boundary clocks.

The set of devices that can communicate using the PTP protocol IP multicast transmissions are said to be in the PTP *subdomain*. This is usually a set of PTP devices connected by a switched network or direct links. The “best” clock in the subdomain is known as the *master clock*. The master clock of a boundary clock is known as the *grandmaster clock*.

The protocol negotiates among PTP ports to identify the device with the highest quality clock source. Ports issuing messages from the *master clock* are said to be *masters*, while those that will receive the messages are *slaves*. When a port will not participate in the protocol its status is *passive*. When the network architect knows the relative quality their clock's time sources, they may configure a specific clock to be the *preferred master*.

PTP Master Election

PTP clocks exchange *SYNC* messages containing information which is used by the PTP *Best Master Clock* (BMC) algorithm. Several factors will affect the choice of best master clock, including the preferred master clock setting, the clock identifier, grandmaster settings and clock stability.

The *clock identifier* is the measure of PTP clock quality and is one of the following:

PTP Identifier	Descrption
GPS	The PTP clock is a primary reference standard traceable to a recognized standard source of time such as GPS. The router uses this identifier when GPS is locked.

NTP	The PTP clock is a secondary reference standard reference clock. The router uses this identifier when it has synchronized with remote NTP server.
DFLT	After the router has power cycled but before any GPS or NTP locks have occurred.

PTP favors preferred masters over normal masters, GPS over NTP over DFLT, higher clock stability over lower clock stability.

Synchronizing NTP from IEEE1588

If GPS is unavailable and PTP becomes a slave the NTP server will view the received IEEE1588 time as any other source of time. The quality (i.e. stratum) of IEEE1588 information is determined by the type of clock source at the master, the number of Boundary Clock hops and the measured network jitter.

The number of Boundary Clock hops is the number of IEEE1588 devices the original time source is relayed through (and not Ethernet hops) and is always 1 or higher.

The measured network jitter factor is 0 if jitter is higher than 1 microsecond and -1 if less than 1 microsecond.

PTP Identifier	Stratum reported to NTP
GPS	1 + Number of Hops ? 1 (if low jitter)
NTP	user configurable value (default 2) + Number of Hops ? 1 (if low jitter)
DFLT	user configurable value (default 10) + Number of Hops ? 1 (if low jitter)

The stratum number reported will be limited to a range of 1 to 16 to comply with NTP.

As an example, a directly connected PTP clock having a GPS clock source and low jitter would report a stratum of 1. With defaults a 2 hop away PTP clock having a NTP clock source and high jitter would report a stratum of 4.

IRIGB Fundamentals

IRIGB outputs are provided by the Precision Time Protocol Card option.

The Inter-Range Instrumentation Group (IRIG) IRIG-B standard details the format of an output signal containing information for the current day, hour, minute and second in UTC format, broadcast at the start of each second. The RuggedRouter complies to IRIG Standard 200-04 generating formats IRIGB002 and IRIGB003 (PWM) and IRIGB122 and IRIGB123 (AM).

IRIGB Output Formats

The router provides three ports by which the signal is distributed, namely:

- An Amplitude Modulated (AM) sinusoidal output port (PTP1),
- Two TTL voltage level output ports (PTP2 and PTP3) which may be configured as either pulse per second (PPS) or pulse width modulated (PWM).

The signal can be used to synchronize intelligent devices to a high quality time source, called the reference clock. The router uses a global positioning satellite (GPS) receiver, NTP or the router's local clock as the reference clock.

Reference Clocks

The GPS provides the highest quality reference clock. It will always be used when it is available, but may require some time after boot before becoming acquired (or "GPS locked"). Typically, GPS lock is usually acquired within five minutes of boot. When GPS is the reference clock, IRIG-B timestamps are accurate to within ns.

If GPS has not yet locked and IEEE1588 is locked, the router will use IEEE1588 server as a reference clock. When IEEE1588 is synchronized, IRIG-B timestamps are accurate to within microsecond or sub microseconds.

If GPS and IEEE1588 have not yet locked, the router will use an NTP server or peer as a reference clock. NTP typically requires less than two minutes after boot to synchronize. When NTP is the reference clock, IRIG-B timestamps can be accurate to within ms.

Before NTP is able to synchronize, the router will use the local clock to obtain the time and will emit IRIG-B timestamps on a one second basis.

How The Router Selects A Reference Clock

The router can be configured to use the following as reference clocks:

- GPS, IEEE1588, NTP and the local clock,
- GPS, NTP and the local clock,
- GPS and IEEE1588,
- GPS

If the router is configured to use multiple reference clocks, it will start sending timestamps using the best ever locked reference clock (local clock is always locked). If better reference clock is locked later, the router will "step" (i.e. suddenly change) the time and use the new reference clock. If the current reference clock becomes unavailable, the router will keep running with its own high precision timing hardware. It will use this hardware until the last used reference clock is locked or a higher quality reference clock is available.

If the router is configured to use only GPS, no timestamps will be issued until GPS locks. If GPS fails, the router will keep running with its own high precision timing hardware. When GPS returns, the time will be stepped back to the GPS reference clock.

GPS Cable compensation

GPS signals received by the antenna will be delayed in time depending upon the type and length of the cable to the router. This delay will introduce inaccuracy in the calculated time and position.

The RuggedRouter provides a method to account for this delay. The table below gives some examples of the delay that can be expected for a given dielectric type. Please note that cable characteristics varies from one manufacturer to the other.

Dielectric Type	**Time Delay in ns/m (ns/ft)**
Solid Polyethylene	4.62 (1.54)
Foam Polyethylene (FE)	3.81 (1.27)
Foam Polystyrene (FS)	3.36 (1.12)
Air Space Polyethylene (ASP)	3.45-3.63 (1.15-1.21)
Solid Teflon (ST)	4.38 (1.46)
Air Space Teflon (AST)	3.39-3.60 (1.13-1.20)

IRIGB/IEEE1588 Main Menu

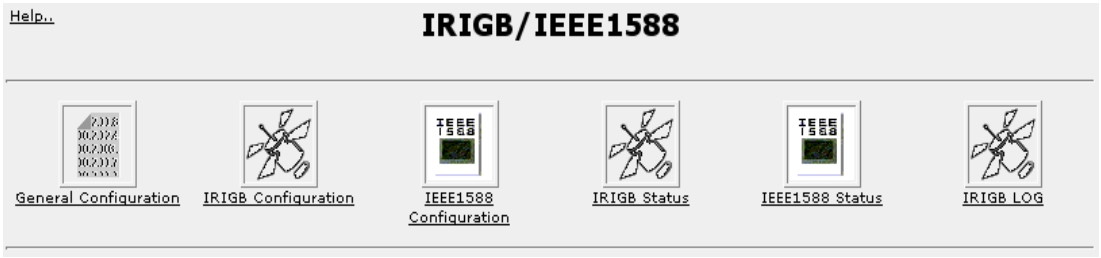

Figure 187: IRIGB/1588 Main Menu

This menu allows you to configure IRIGB and IEEE1588, display its current status and review historical changes.

General Configuration

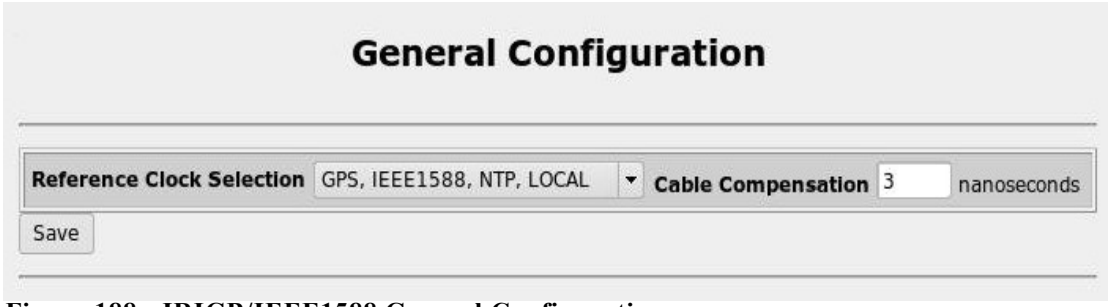

Figure 188: IRIGB/IEEE1588 General Configuration menu

This menu allow you to configure general parameters.

The **Reference Clock Selection** field selects the order in which to prefer reference clocks.

The **Cable Compensation** field specifies the value, in nanoseconds, that will be used to compensate for the cable type and length. The compensation is done using integer nanosecond values. Fractional decimal values will be truncated.

IRIGB Configuration

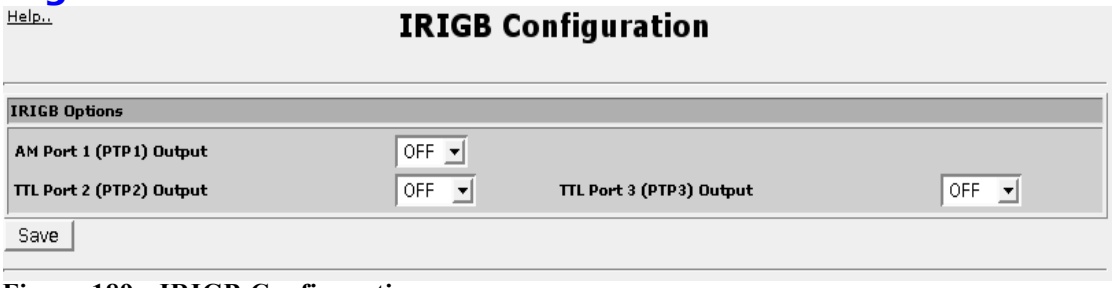

Figure 189: IRIGB Configuration menu

This menu allow you to configure IRIGB parameters. The save button will save the changes of configuration permanently.

The **AM Port 1 (PTP1) Output** field enables or disables the amplitude modulated output of this port.

The **TTL Port 2 (PTP2) Output** and **TTL Port 3 (PTP3) Output** fields sets the output formats of these ports to PPS, PWM and OFF.

IEEE1588 Configuration

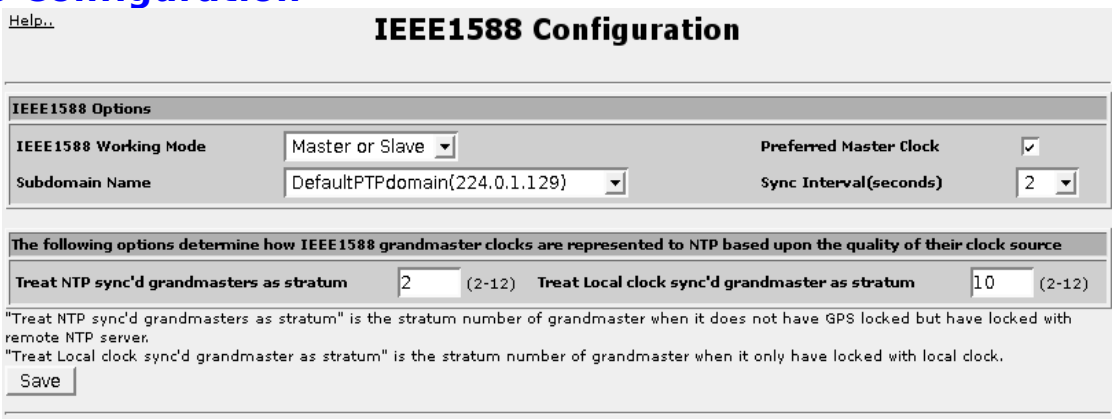

Figure 190: IEEE1588 Configuration Menu

This menu allows you to configure IEEE 1588 parameters.

The **1588 Working Mode** field allows configures whether the router will be forced to 1588 slave mode or determine its role by the BMC algorithm.

The **Preferred Master Clock** field configures the router to be preferred master clock.

The **Subdomain Name** field allows you to choose which domain you want the router to participate in. There are four domains available, each mapped to a different multicast IP address.

The **Sync Interval** field configures the rate at which SYNC messages are issued.

The router NTP daemon uses GPS as a clock source when it is available and with IEE1588 when GPS is not available.

The **Treat NTP sync'd grandmaster as stratum** field assigns the stratum number when grandmaster clock synchronized with remote NTP server but not GPS.

The **Treat Local Clock sync'd grandmaster as stratum** field assigns the stratum number when grandmaster clock synchronized with local clock but not NTP server or GPS.

IRIGB Status

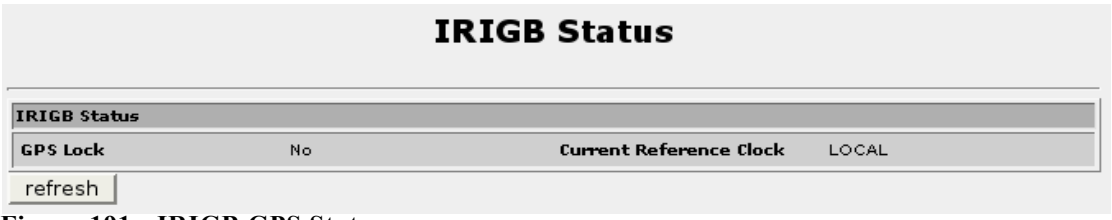

Figure 191: IRIGB GPS Status

This page shows whether GPS is locked, and the source of the current reference clock.

IEEE1588 Status

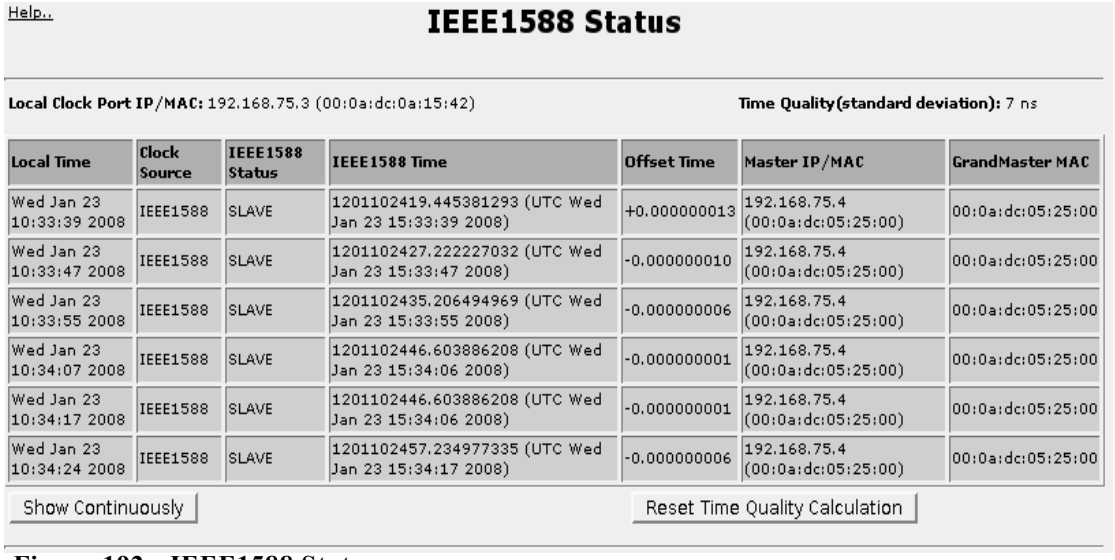

Figure 192: IEEE1588 Status

This page shows the historical status of IEEE1588 on the router.

The line above the table provides the local clock IP address, MAC address and the time quality information. The table will contain entries made when the clock source or status changes. The current local time on the router, the IEEE1588 status, IEEE1588 and UTC time, the offset from master in seconds, the master IP/MAC address and grandmaster MAC address are provided.

IRIGB Log

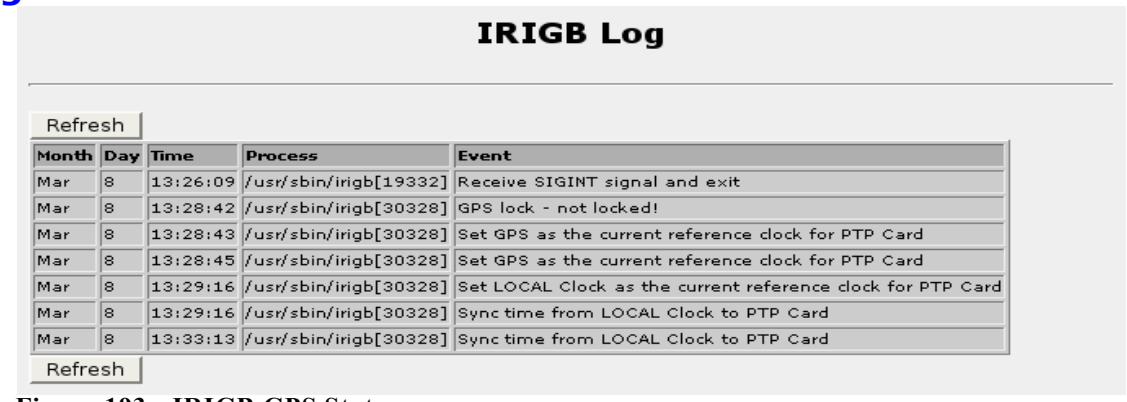

Figure 193: IRIGB GPS Status

This page reflects reference clock changes in IRIG-B.

This page intentionally blank

Chapter 25 – Configuring The Snort IDS

Introduction

This chapter familiarizes the user with:

- Configuration of Snort as an Intrusion Detection System.
- Generating a daily snort analysis email.

Snort Fundamentals

The snort Intrusion Detection System (IDS) provides a type of security management system for the router. Snort gathers and analyzes information on various network interfaces to identify possible security breaches, which include both intrusions (attacks from outside the protected network) and misuse (attacks from within the protected network).

Snort examines packets received on selected interfaces, applies "rules" from its database and generates "alerts" to warn of "vulnerabilities".

Which Interfaces To Monitor

Typically, the router will have an interface to an external network and interfaces comprising the local network. The firewall will cite these interfaces as belonging to the net and local zones. A key decision is whether to monitor traffic outside, or inside of the firewall.

Monitoring traffic outside the firewall (on the external network interface) has the advantage that attacks the firewall is blocking can be seen. This method, however, will generate a large number of alerts. Additionally, firewall rules installed to eliminate vulnerabilities will not prevent future alerts since traffic is monitored *before* the firewall. Finally, this method will not detect misuse of the local ports.

Monitoring traffic inside the firewall (on all local interfaces) has the advantage that the number of alerts decreases as vulnerabilities are eliminated at the firewall. It's also good to monitor as much of the internal traffic as possible.

Snort Rules

The router supplies a variety of prepackaged rules. Each rule contains a unique Signature Identifier (SID). The SID is included in reported alerts as part of a Snort unique rule ID, a three digit number of the form [generator:SID:revision]. The "generator" field reflects the organization that generated the rule, official snort rules having values less than 1,000,000. The SID is a unique number to reflect an individual rule, while the "revision" reflects improvements to the rule.

The main Snort IDS menu provides the capability to disable individual and groups of rules. It is also possible to add unique rules to the database and to replace the existing set of rules with more experimental rules from the community.

Alerting Methods

Alerts generated by snort are stored by one of three methods; as local syslog messages, remotely sylogged messages and in an alert file.

When the local syslog method is chosen, the destination log file may be selected.

When the alert file method is chosen, a daily analysis of the file can be emailed.

The SIDs referenced in alerts can be used to quickly locate the rule via the main Sort IDS menu. The rule itself often contains HTML links to Internet resources such as www.securityfocus.com and cve.mitre.org. These provide more in depth descriptions of the vulnerability.

Performance And Resources

The performance impact of snort varies with the number of interfaces monitored, the number of rules enabled, the packet rate and the logging method.

Snort has been empirically determined to use about 20% of the CPU clock cycles at its maximum processing rate.

The router is capable of recording about 300 entries/second to the local syslog and 500 entries/second to the alert file. Alerts at rates exceeding the above rates will not be recorded.

Snort will require 5 Mbytes of system memory to start with an additional 15 Mbytes of memory for each interface monitored.

Snort IDS Main Menu

This menu configures the snort IDS and is composed of three sections.

Note that snort is disabled by default and may be enabled via the System folder, Bootup And Shutdown menu. If snort is running, configuration changes must be made active by restarting it. The Restart Snort button will restart snort, listing the interfaces it is active upon.

Global Configuration

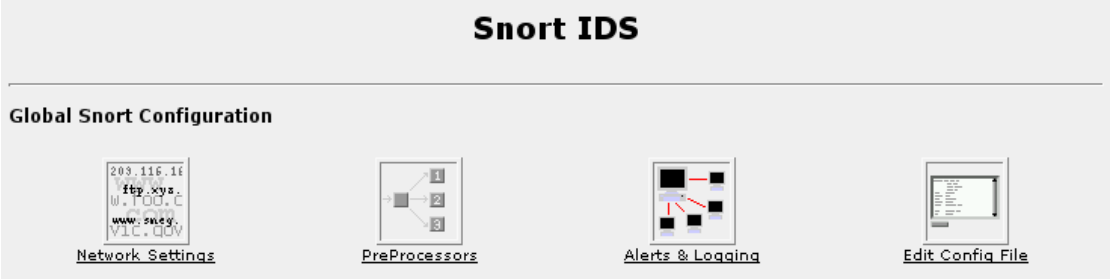

Figure 194: Snort Main Menu part 1

The Global Configuration menu section configures parameters that apply to all interfaces.

Interfaces

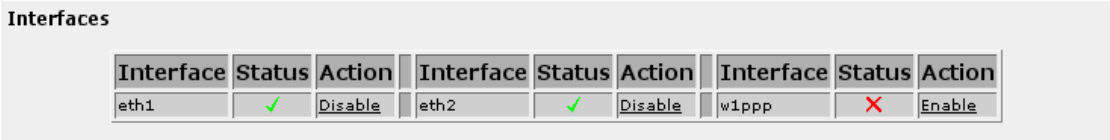

Figure 195: Snort Main Menu part 2

The Interfaces section selects the interfaces snort will monitor. You must restart snort after changing interfaces.

Rulesets

Figure 196: Snort Main Menu part 3

The Rulesets section selects the rules to apply on monitored interfaces.

Each "ruleset" reflects a collection of rules that are related. The link under the **Action** field will disable or enable all of the rules in a ruleset. Individual rules in a ruleset may be modified by following the set name link under the **Rule Set** field, resulting in a menu such as the following.

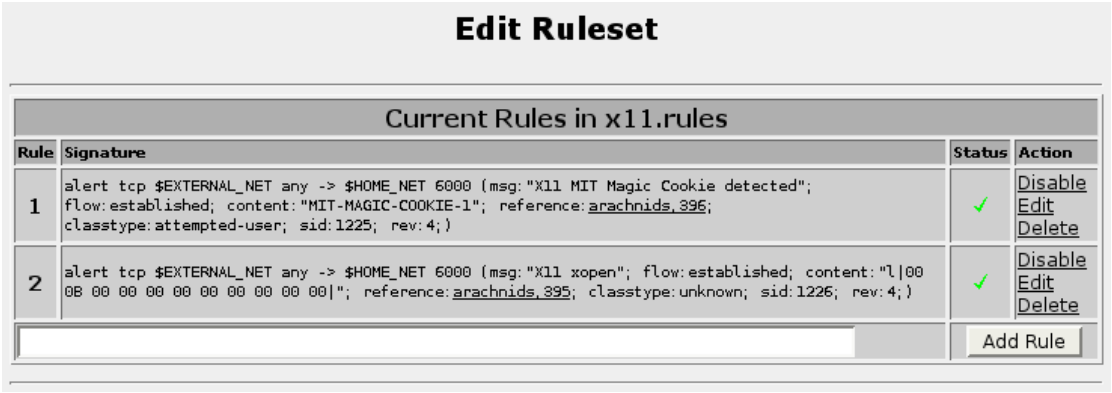

Figure 197: Snort Ruleset Edit

Each rule can be individually enabled, disabled or deleted. Most rules will include a reference link to more information about the vulnerability the rule detects.

It is possible to add your own rule, or one obtained from the open source community (e.g. www.bleedingsnort.com).

Rule Lookup by SID

The Look Up Rule button accepts a SID and displays its rule. You may elect to disable the rule or learn more information about it.

Network Settings

Figure 198: Snort Network Settings

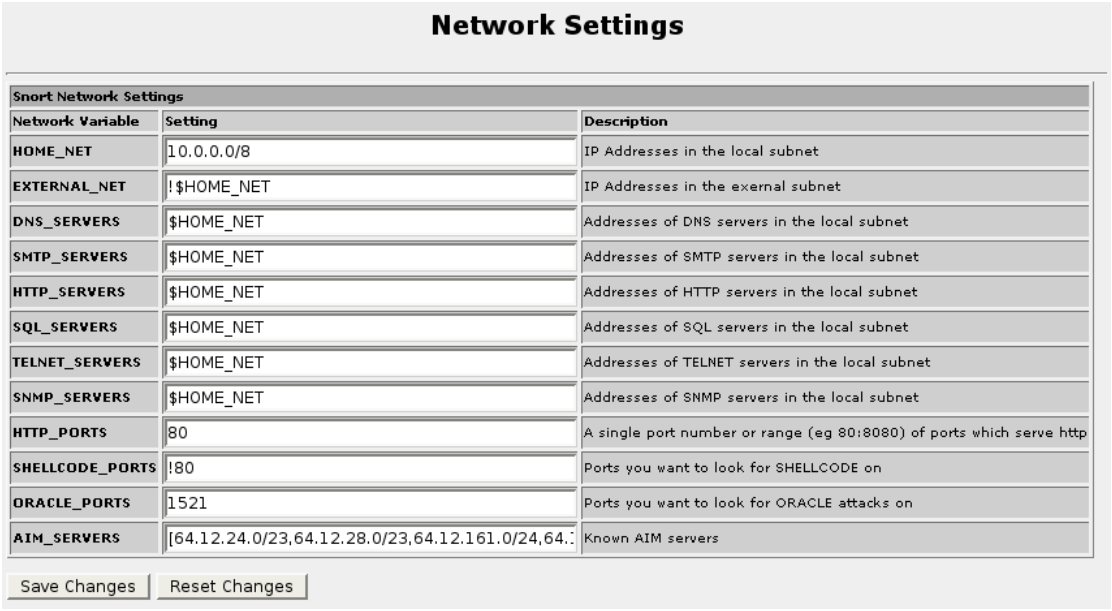

This menu allows you to configure the IP addresses and ports of servers in the local and external network.

The **Home Net** field defaults to “ANY” and designates the IP subnet of any local ports on the router. Configuring a specific subnet can reduce the number of alerts generated.

PreProcessors

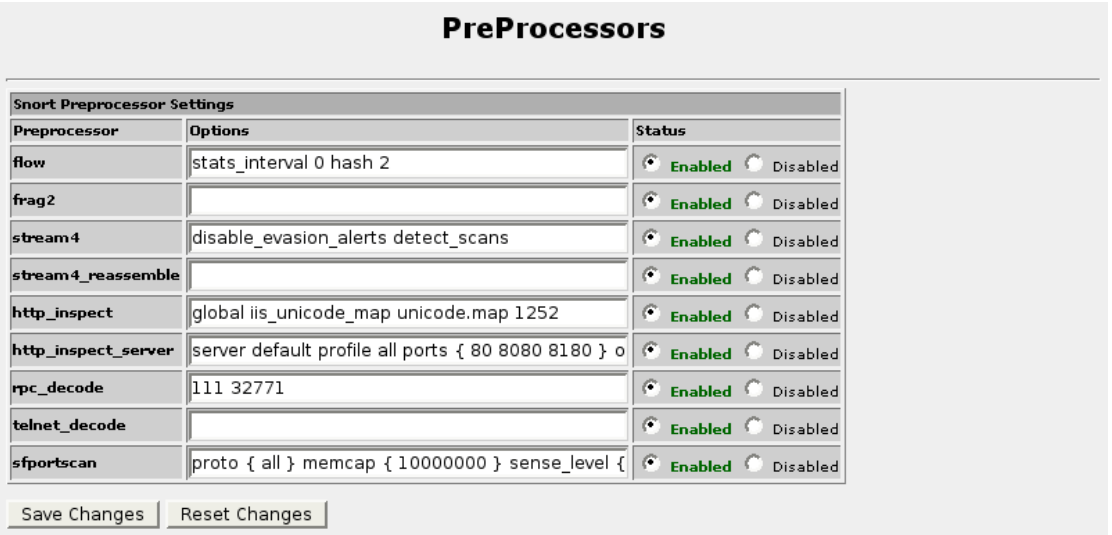

Figure 199: Snort Preprocessors

Preprocessors are plug-in modules that operate on the captured packets. Preprocessors perform a variety of transformations to make it easier for snort to classify packets.

The configuration of preprocessors is beyond the scope of this user guide.

Alerts & Logging

Figure 200: Snort Alerts

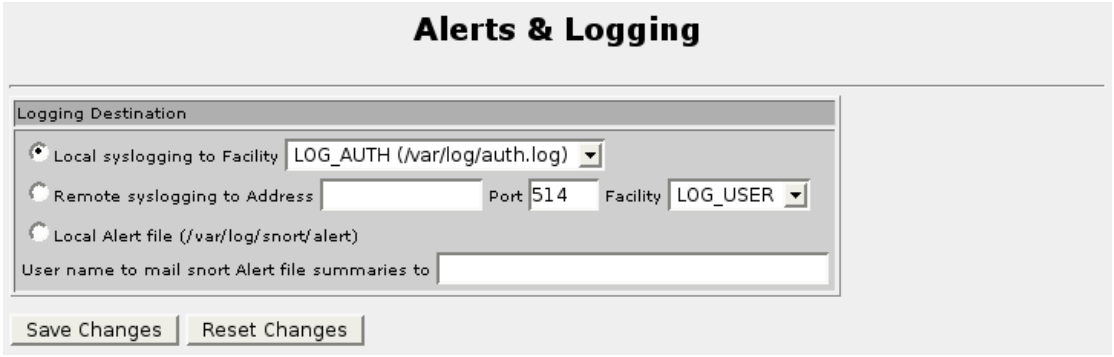

Alerts generated by snort are stored by one of three methods; as local syslog messages, remotely sylogged messages and in an alert file.

When the **Local syslogging** method is chosen, the destination log file may be selected.

When the **Remote syslogging** method is chosen, the IP address of the remote syslog host must be identified. Modifying the **Facility** field will determine how the alert is logged on the remote host.

When the alert file method is chosen, a daily analysis of the file can be emailed to the user provided in the **User Name..** field. Note the you must also visit the Maintenance menu, Miscellaneous sub-menu, Outgoing Mail sub-menu in order to configure a mail forwarder.

Edit Config File

Snort is extremely flexible and not all capabilities have been described in this user guide. This menu provides the user with the ability to make raw configuration changes to the snort configuration file from within Webmin.

Chapter 26 – Maintaining The Router

Introduction

This chapter familiarizes the user with:

- Viewing Alerts
- Configuring and monitoring the Gauntlet Security Appliance
- Backing up and restoring configurations
- Configuring SNMP
- Configuring Radius Authentication
- Configuring Outgoing Mail
- Using System Logs
- Upgrading Software
- Using Pre-upgrade/Post-upgrade scripts
- Uploading and downloading files

Alert System

The alert system provides the following features:

- Generates alerts, displaying them locally and/or forward them via email messages.
- Alerts are set and cleared by the daemons that own them. Active alerts are locally displayed and can be cleared manually.
- Multiple forwarders can be configured, a configurable filter level controls alert forwarding to each forwarder.
- By configuring different forwarders, low severity and high severity control centers can be set up.

Each alert is mapped to an alert definition entry, which is predefined by a daemon who owns the alert or by a user. All alert definition entries are configurable by user.

An alert filter is a user defined configuration to define the forwarders destination of active alerts. Any active alerts with Renotify Interval set to non-zero value and matches with the filter level will be forwarded to the defined forwarder destination.

Alert Menu

Figure 201: Alert Main Menu

This menu displays active alerts and allows you to change alert system configuration and alert definitions.

Follow the **All Alerts** link to show all alerts. Follow the severity links (**Emergency .. Debug)** or the category links **(chassis .. daemon)** to to limit the alert view.

Note that active alerts are volatile and will be regenerated after reboot. If you clear an alert manually, it will appear if the condition occurs again. You may disable the alert permanently by disabling the alert from its entry in the definition menu.

The **Clear Alert** link under the **Action** column allows you to clear the alert.

Clicking on the Alert Name, Specific, Severity and Date column headers will sort the alerts by those types.

Select **Alert Configuration** to change the generic configuration and alert filter configurations.

Select **Alert Definition configuration** to change the alert definition entries.

Alert Configuration

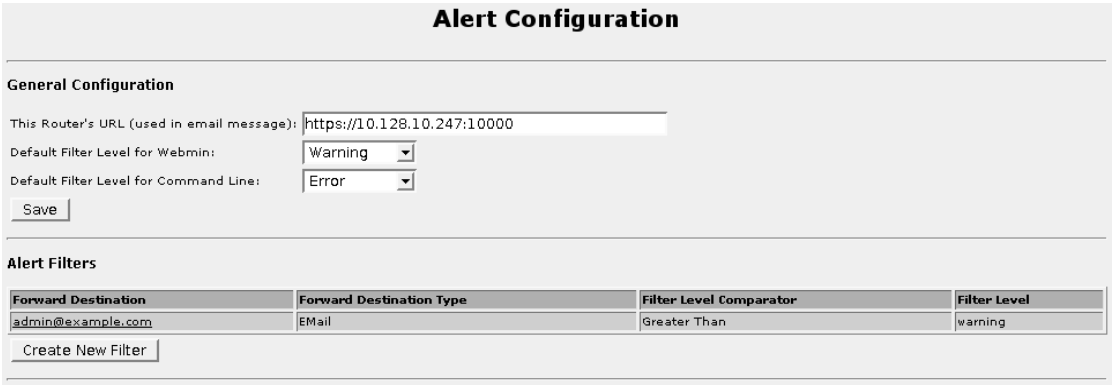

Figure 202: Alert Configuration Menu

This menu configures the general information and forward filters for the alert system.

The **This Router's URL** configures the link to access this router. This information will be used in the email forwarder, which user can click on the link in the email to access the router.

The **Default Filter Level for Webmin** configures the lowest alert level to show on webmin. All active alerts higher priority than this level will be displayed on the webmin home page.

The **Default Filter Level for Command Line** configures the lowest alert level to show when user login by console or ssh.

The **Save** button saves all changes of general configuration.

The **Create New Filter** button allows you to create a new forwarder filter for active alerts.

Alert Filter Configuration

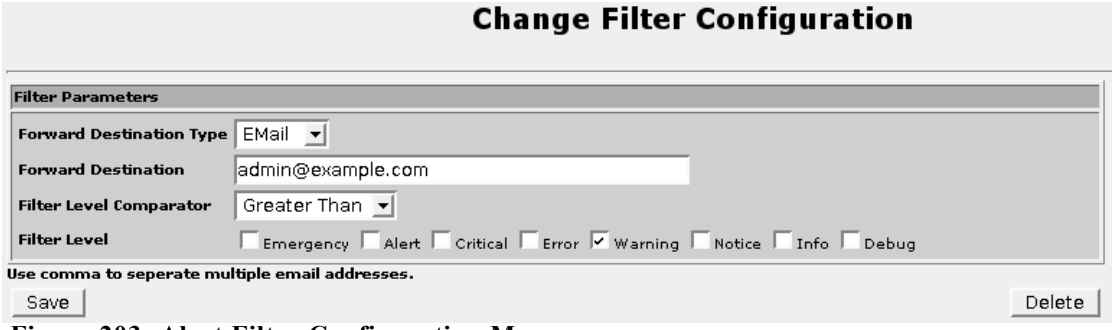

Figure 203: Alert Filter Configuration Menu

This menu configures an alert filter, which defines the forwarder destination for active alerts matching with defined filter level.

The **Forward Destination Type** configures the type of filter. Currently only type Email is supported.

The **Forward Destination** configures the destination matching with the Forwarder Destination Type. Note that multiple email addresses should be separated by comma.

The **Filter Level Comparator** configures the way to match with defined filter level.

The **Filter Level** configures what filter level is to be compared. Note that Emergency has the greatest filter level and Debug has the lowest filter level.

Alert Definition Configuration

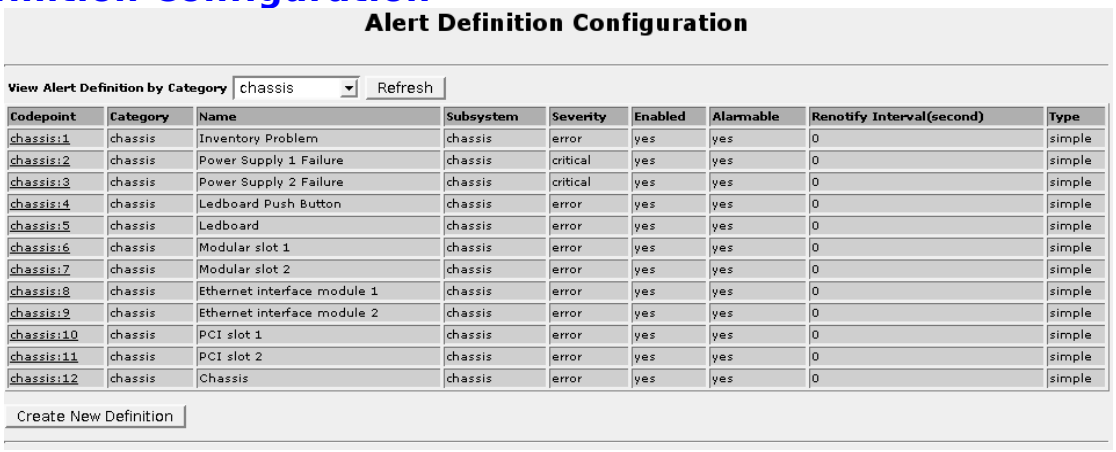

Figure 204: Alert Definition Configuration Menu

This menu displays matched alert definition entries. It also allows user to change the an alert definition entry or create a new entry.

An alert definition entry defines an alert which will be monitored by the system.

The **View Alert Definition by Category** allows you to display alert definition entries matching with selected category.

The **Create New Definition** button allows you to create a user defined alert definition entry.

Click on one of the link under the **Codepoint** column allows you to change the configuration for that alert definition entry.

Change Alert Definition

Figure 205: Change Alert Definition Menu

This menu allows you to change an existing alert definition entry.

The **Codepoint** is the key part of the alert definition entry and does not allow to be changed.

The **Category** configures which category the alert definition entry belongs to.

The **Name** configures the name of the alert definition which will be displayed by webmin, login or email forwarder when an active alert exists.

The **Subsystem** configures which subsystem the alert definition entry belongs to.

The **Severity** configures the severity level of the alert. The severity level is sorted from highest priority to lowest priority.

The **Alarmable** configures whether the matched alert should trigger the critical relay and alarm LED on the LED panel of the router.

The **Enabled** configures whether the alert system should monitor and record matched active alert. If Enabled is not checked, matching active alert will be ignored.

The **Renotify Interval** configures how often should the matched active alert be notified according to alert filter configuration setting. If it is disabled, no notification will be forwarded.

The **Type** configures type of the alert definition entry. There are three types available: **Simple**, **Shell** and **RMON**. Currently only the first two types are supported. If users choose **Shell** type, they should complete parameters under **Parameters for Shell** table.

The **Parameters for Shell** table allows user to configure additional parameters if the alert definition entry type is **Shell**.

The **Sample Interval** configures how often should the system run configured shell command to get a sample.

The **Command** configures the shell command to run.

The **Comparator** configures how to compare with the shell command result.

The **Threshold** configures the threshold to compare with the shell command result to see whether the condition is true or false.

The **And Repeats** configures how many times the condition must be true before the alert is generated.

The **And Until** configures how many seconds the condition should be true before an alert is generated.

The **Not Cleared Repeats** configures how many times the condition must be false before the alert is cleared.

The **Not cleared Until** configures how many seconds the condition must be false before an alert is cleared.

The **Parameters for RMON table** allows user to configure additional parameters if the alert definition entry type is RMON type.

The **Device Name** configures the name of the device to be monitored.

The **MIB Variable** configures the MIB variable being monitored.

The **Sample Interval** configures how often samples should be generated.

The **Rising Threshold** configures the value that will trigger an event when the value of the variable increments past this value.

The **Falling Threshold** configures the value that will trigger an event when the value of the variable decreases past this value.

The **Startup** configures the condition that will cause the initial event.

Gauntlet Security

RX1100 owners can use the Gauntlet security appliance to restrict access to critical assets. This section details how to activate Gauntlet and determine currently negotiated sessions. Details and recommendations on applying the Gauntlet system to networking may be found in texts referenced in the **About This Guide** section of the user guide.

What And How Gauntlet Protects

Gauntlet protects against unauthorized access to critical assets, including the router itself. Gauntlet allows connection from known management devices to assets behind the firewall operating on known TCP/UDP port numbers. Gauntlet does not encrypt communications which occur in the clear, such as the Telnet protocol. Protocols such as SSH and HTTPS offer their own encryption and are suitable for use with Gauntlet.

Gauntlet And The Firewall

Gauntlet integrates tightly with the firewall, opening it for communications between vetted clients and critical assets on a demand basis. There are three steps in activating the Gauntlet security appliance.

1. The firewall must be configured with some default rules required by the appliance (described below), and then activated or restarted.
2. The rrsetup utility must be used to configure a Gauntlet passphrase and enable the Gauntlet daemon.
3. The Gauntlet daemon and Shorewall must both be enabled in the Webmin Bootup and Shutdown Menu to "Start at boot".

Shorewall requires you to assign the router interfaces to zones and then control traffic between these zones.

Typically, the zone for WAN interfaces is named "net" while the zone for local interfaces is named "loc". The following instructions assume those names.

The gauntlet daemon requires rules for certain ports (shown below) to be installed. Contact RuggedCom support for assistance if you wish to reassign these ports.

1. Visit the Shorewall Network Zones sub-menu and create the net and loc IPv4 zones.
2. Visit the Network Interfaces sub-menu and assign interfaces to the zones.
3. Visit the Default Policies sub-menu and assign the following policies:

Source zone	Destination zone	Policy
fw	any	ACCEPT
loc	net	ACCEPT
all	any	REJECT

4. Visit the Firewall Rules sub-menu and assign the following rules:

Action	Source zone	Destination zone	Protocol	Src-Port	Dst-Port
ACCEPT	net	fw	UDP	any	30000
ACCEPT	net	fw	UDP	any	30001
Gauntlet	net	loc			
Gauntlet	net	fw	TCP	any	31000
Gauntlet	net	fw	TCP	any	31002

```
Gauntlet    net         fw          TCP     any     10000
```

The order of rules is significant. Rules inserted before this set will not be protected by Gauntlet. Any rule appearing after the gauntlet chain rules will automatically be ignored. Consult with RuggedCom support for assistance. If you want to grant SSH access to the router, replace "10000" in the last rule with "22,10000".

When adding these rules via Webmin, for those rules where you select "Gauntlet" from the "Action" pulldown list, be sure to leave the "log to syslog level" set to "<Don't log>". If you manually edit the "/etc/shorewall/rules" file then do not specify any loglevel in your Gauntlet rules.

5. Ensure that the firewall is enabled in the Bootup and Shutdown Menu and apply the firewall configuration to effect the changes.

Note: You must ensure that the firewall is configured and enabled when using the Gauntlet Security Appliance.

Gauntlet Status Menu

Gauntlet integrates tightly with the firewall, opening it for communications between vetted clients and critical assets on a demand basis.

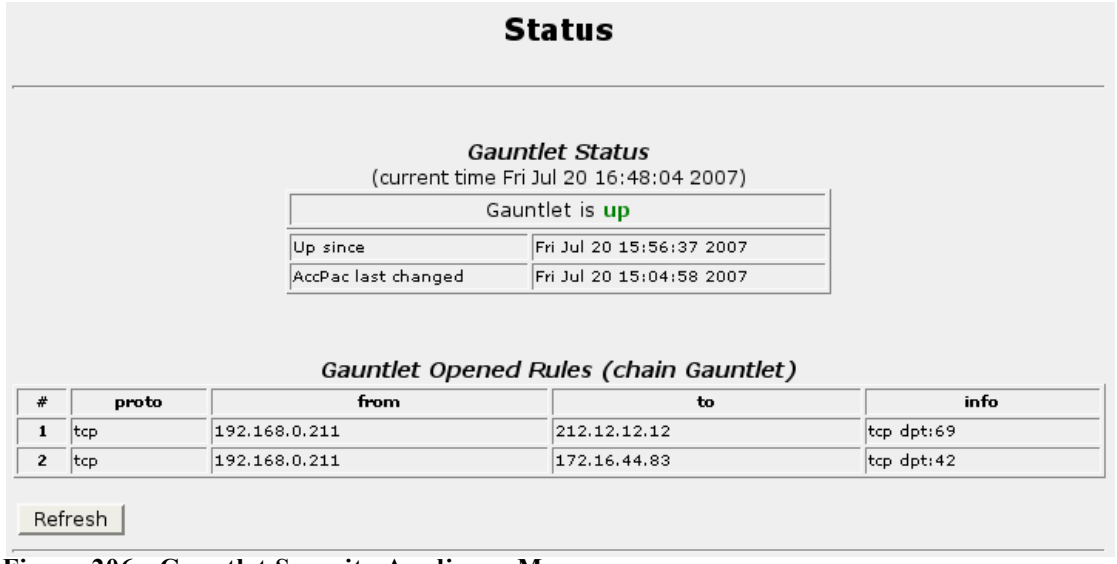

Figure 206: Gauntlet Security Appliance Menu

The status menu provides a list of validated open connections.

Upgrading Gauntlet

During an upgrade, the Gauntlet daemon may be required to restart. During the upgrade all existing Gauntlet protected connections will be closed.

Backup And Restore

Figure 207: System Backup And Restore

The Backup And Restore system provides the following features:

- All configuration settings are saved in a configuration archive,
- Archives can be used to "clone" routers, replicate a damaged resource or unwind a change,
- Archives can be created manually (including user comments) or by the Automatic nightly backup, which captures all changes over the previous 24 hours,
- The nightly backup archives can be automatically transferred via scp or ftp to a designated server,
- The nightly backup archives are kept on the router for a configurable number of days and then deleted. The most recently made archive is never destroyed. Manually created archives are never destroyed.
- If you make a configuration change you later wish to reverse you can restore a previously made archive completely. An archive difference tool is provided, showing the difference between one archive and either another archive or the current configuration. Changes in configuration can also be detected and "unwound" by applying the previous state of a router on a file by file basis.
- Archive filename is user definable and can include any of date/time, host name and/or release version,
- Archives can be uploaded to the router and restored. The router prevents the restoration of archives having other than current software version.
- A factory defaults file is included.

Note the following caveats:

- Chassis specific items such as serial number, hardware inventory and MAC addresses are not saved,
- Log and history files are not saved,
- Information stored in the root and user accounts are not saved.

General Configuration

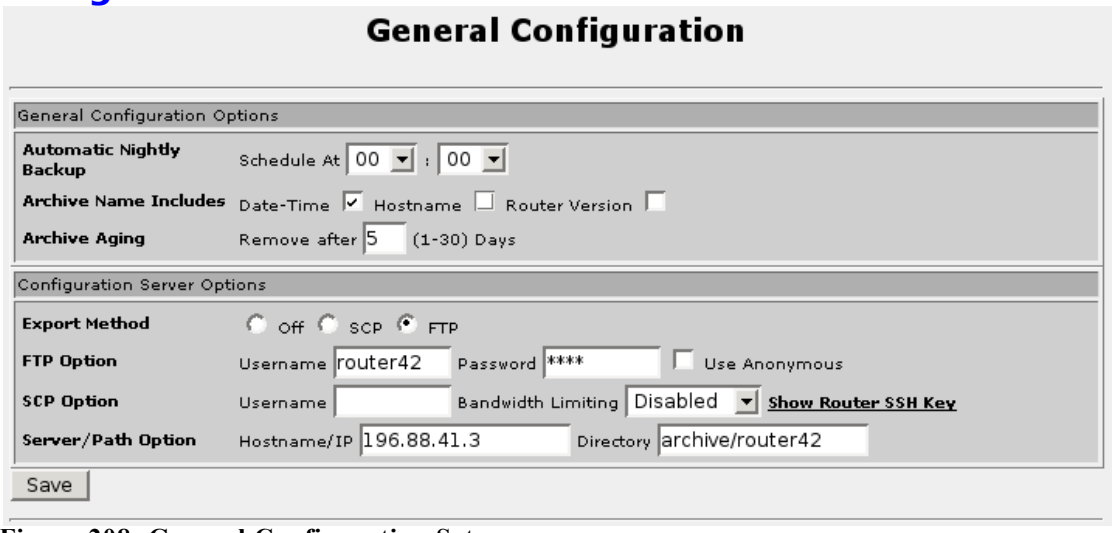

Figure 208: General Configuration Setup

This menu configures the backup system.

The **Automatic Nightly Backup** field specifies when the nightly backup is scheduled. The automatic export to a server will start (if enabled) immediately after the backup completes.

The **Archive Name Includes** field selects text fields (Date-Time, Hostname, Router Version) included in archive name.

The **Archive Aging** field specifies how long nightly backup archives are kept. Note that the most recently made nightly backup will never be deleted. Manually made archives are never aged and must be manually deleted.

The Configuration Server Options table allows user define the configuration server.

The **Export Method** field selects the method of exporting backup archives to a server. If the **Export Method** field is set to FTP, the FTP Options are used. If the **Export Method** field is set to SCP, the SCP Options are used.

The **FTP Option** field specifies FTP User name, Password or to use anonymous FTP .

The **SCP Option** field specifies SCP User name and Bandwidth Limitation when the Export Method is SCP. The **Show Router SSH Key** link will display the ssh public key for this router, which can be used in the configuration server to accept SCP from the router.

The **Server/Path Option** field specifies the configuration server hostname (or IP address) and the directory in which to save archives.

Archive History

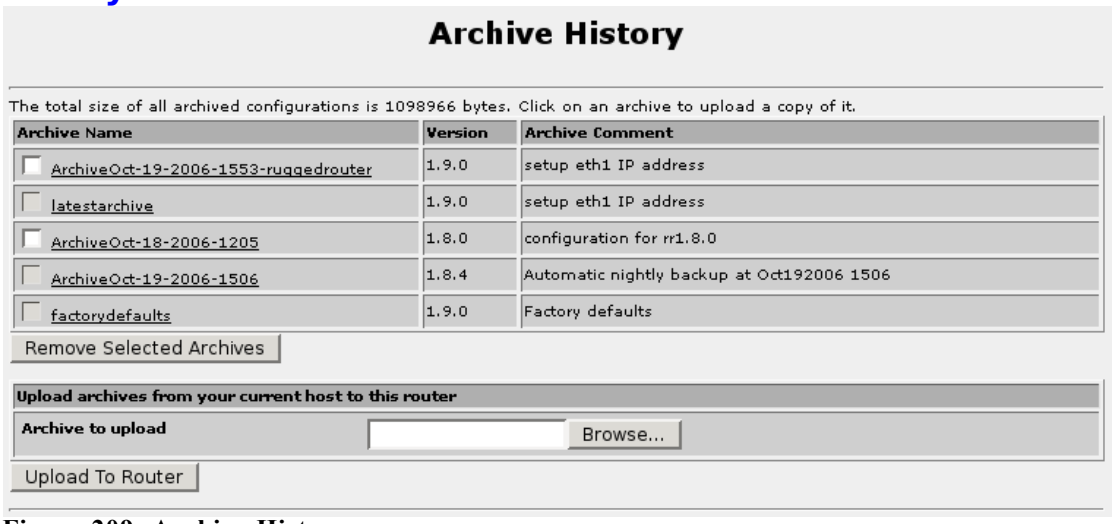

Figure 209: Archive History

The Archive History menu displays current archives, sorted by date (most recent first). Following the link of an archives under the **Archive Name** field upload a copy of it.

Selecting an under the **Archive Name** field and applying the **Remove Selected Archives** button will delete the archive. Note that only manually backup archives can be deleted. Automatic nightly backup archives will automatically aged out . The latestarchive and factorydefaults archives will never be deleted.

The **Archives to upload** fields select archives to upload to the router. The **Browse...** button will allow you to select an archive. Applying the **Upload to Router** button will upload the specified archive to the router.

Archive Backup

Figure 210: Archive Backup

This menu allows the user to manually create an archive. It accepts a comment which will be included in the archive file. The input box above the Start Backup button shows the candidate archive file name, which can be changed by user. Starting the backup results in the following display.

Figure 211: Archive Backup, Complete

The created archive can be immediately uploaded if desired by following the "Upload A Copy Of This Archive.." link.

***Note: If you use the Internet Explorer web browser, you must "Right-click" the link and save the file manually**. Otherwise Internet Explorer will rename the file after uploading, preventing its use in a subsequent archive restore.*

Archive Restore

Figure 212: Archive Restore Menu

The restore process begins by selecting an archive to restore from. Following an archive link will restore the archive and reboot the router.

***Note: Some manually (and even automatically) created archives are not possible to restore.** If the router was upgraded **after** the archive was created, the archive will have old, confusing and possible missing configurations. The **Version** field indicates this. The latestarchive and factorydefault archives always have the current release version (and are always able to be restored). If an archive has a lower version number, it will not be restorable.*

The latestarchive and factorydefault archives are always able to be restored.

Click on one of the links under **Archive Name** to start the restore. Starting the restore results in the following display.

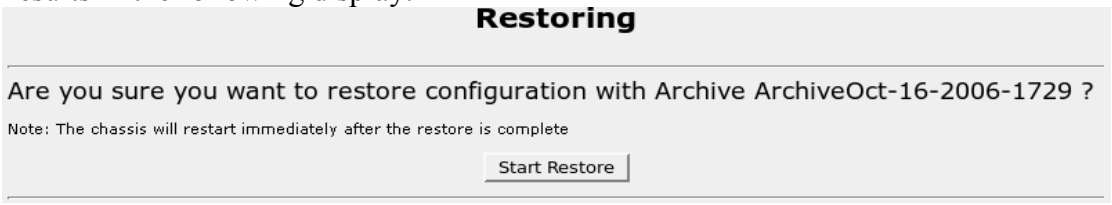

Figure 213: Start Restore

To begin the restoring process, click the **Start Restore** button.

Archive Difference Tool

Figure 214: Archive Differences Menu

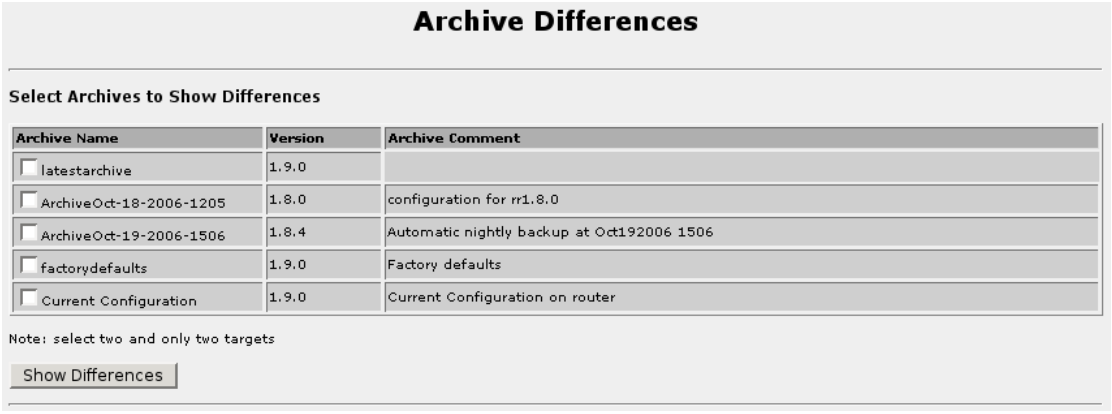

The Archive Difference menu shows the difference between two targets. The first target must be an archive while the second target can be either another archive or the current configuration.

Choose two and only two targets and click the **Show Differences** button.

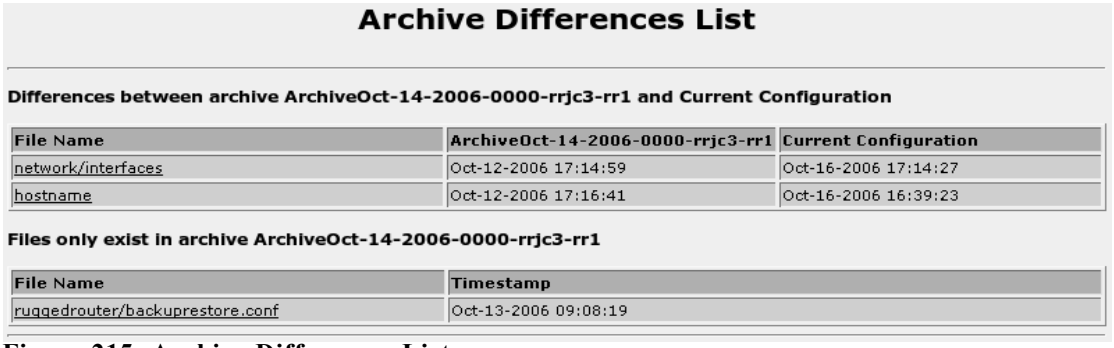

Figure 215: Archive Differences List

The resulting menu shows the differences between the two selected targets. Files in this table are sorted by the change time (most recent changes first). Files that exist in only one of the targets are shown separately.

Following the links under **File Name** column will show a files difference between the two targets.

The difference will be shown by two methods. The difference between the two targets will be first be shown in a side by side scrollable comparison.

The difference will also be shown in a window that shows differing lines.

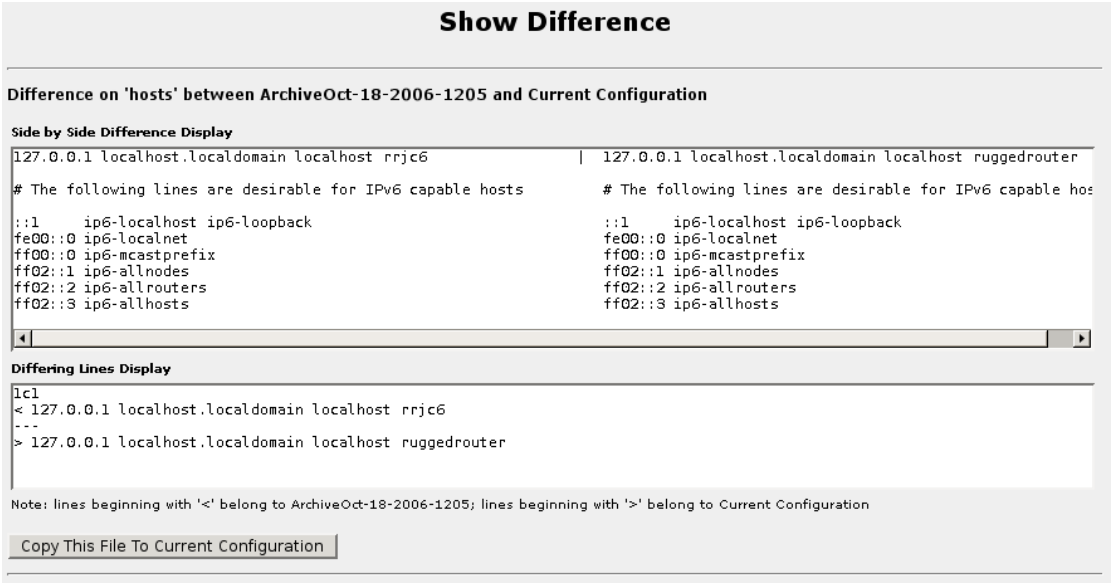

Figure 216: Show Difference for selected file between two targets

The **Copy This File to Current Configuration** button will be present when the destination archive is the Current Configuration. It allows user to copy the selected file from the old archive to current configuration.

Note: It is possible to damage your router through use of this feature! *Ensure that the configuration file copied makes sense in the current version of the router.*

Note that the copying configurations may not make any actual operating changes until the systems that own them are restarted.

If the source archive has a file that is not present in the Current Configuration, it is possible to view that file and then copy it into Current Configuration.

SNMP Configuration

The SNMP menus provide the following configuration features:

- System information
- agent network addresses
- Community access to the agent
- SNMP trap delivery

The SNMP (the Simple Network Management Protocol) protocol is used by network management systems and the devices they manage. SNMP is used to manage items on the device to be managed, as well as by the device itself, to report alarm conditions and other events.

The first version of SNMP, V1, provides the ability to to send a notification of an event via “traps”. Traps are unacknowledged UDP messages and may be lost in transit. SNMP V2 adds the ability to notify via “informs”. Informs simply add acknowledgment to the trap process, resending the trap if it is not acknowledged in a timely fashion.

SNMP V1 and V2 transmit information in clear text (which may or may not be an issue depending the facilities the data is transmitted over) and are lacking in the ability to authenticate a user. SNMP V3 adds strong authentication and encryption.

SNMP Configuration Main Menu

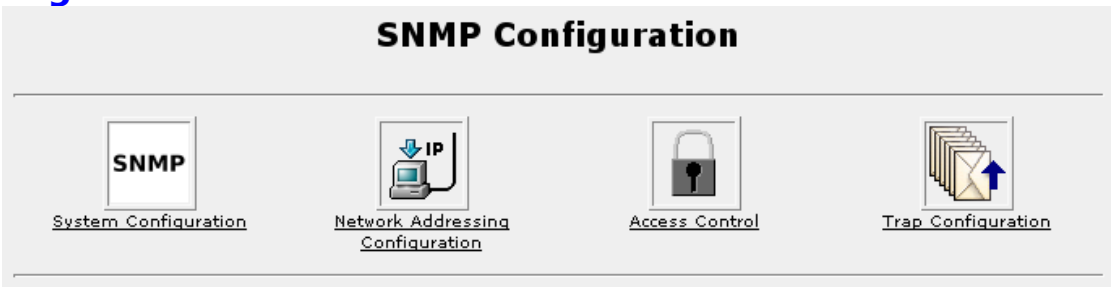

Figure 217: SNMP Main Configuration page

In order to enable snmpd (the snmp daemon) at each and every boot, use the System folder, Bootup And Shutdown menu.

Note: *Prior to ROX 1.10.0, SNMP was manually configured used the com2sec, group, view and access directives. If so configured, the SNMP menu will prompt you to convert the configuration to one it can manage.*

System Configuration

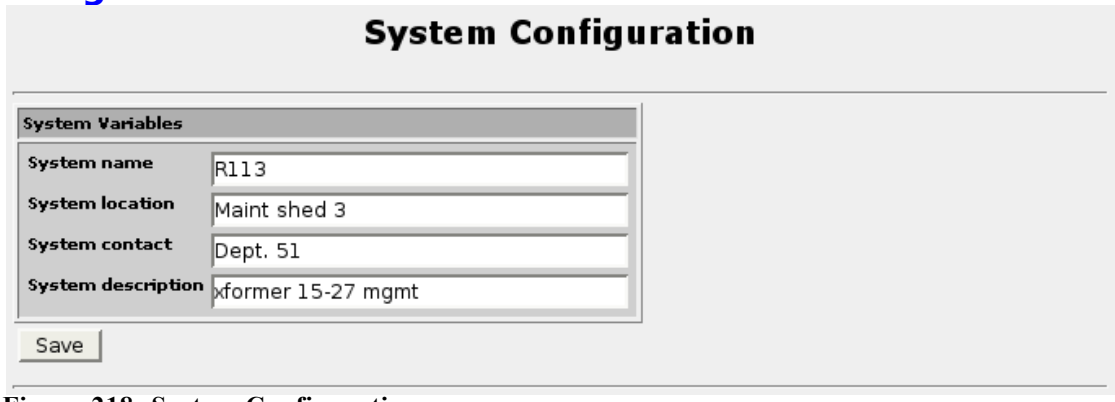

Figure 218: System Configuration page

The **System name**, **System location, System contact,** and **System description** fields configure descriptive parameters for the router.

Network Addressing Configuration

For reference, the set of currently configured and active IP addresses is listed near the the top of the page.

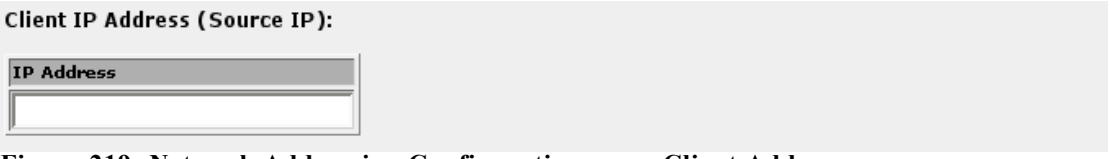

Figure 219: Network Addressing Configuration page, Client Address

The **Client address (Source IP)** field specifies the address from which snmpd will send notifications. If the field is blank, the default behaviour will be to transmit the notification from the IP address of the interface from which the message leaves the router. Snmpd will return to this behavour if the configured address is not available when it starts.

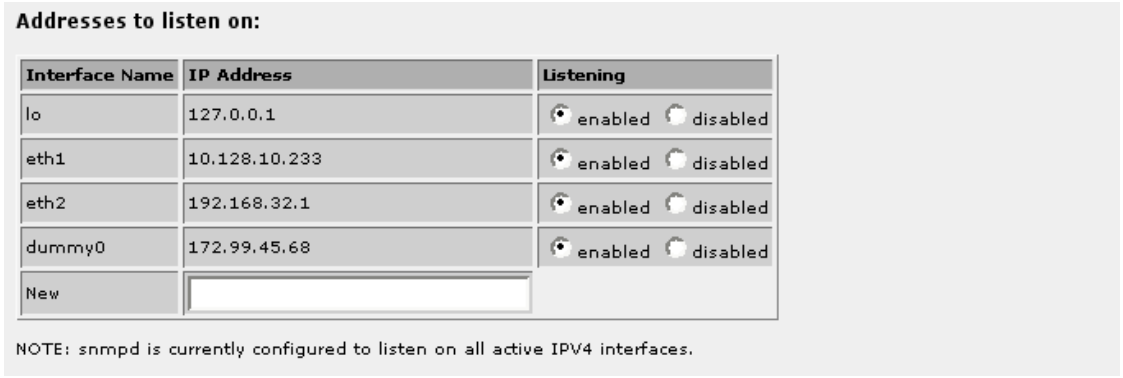

Figure 220: Network Addressing Configuration page, Addresses to listen on

The table of **Addresses to listen on** includes the list of currently configured and active IP addresses, and whether the address is currently listened on. The **New** field allows for the addition of other IP addresses.

Snmpd will use these addresses providing they are active at the time it starts.

By default, snmpd listens on all interfaces.

Access Control

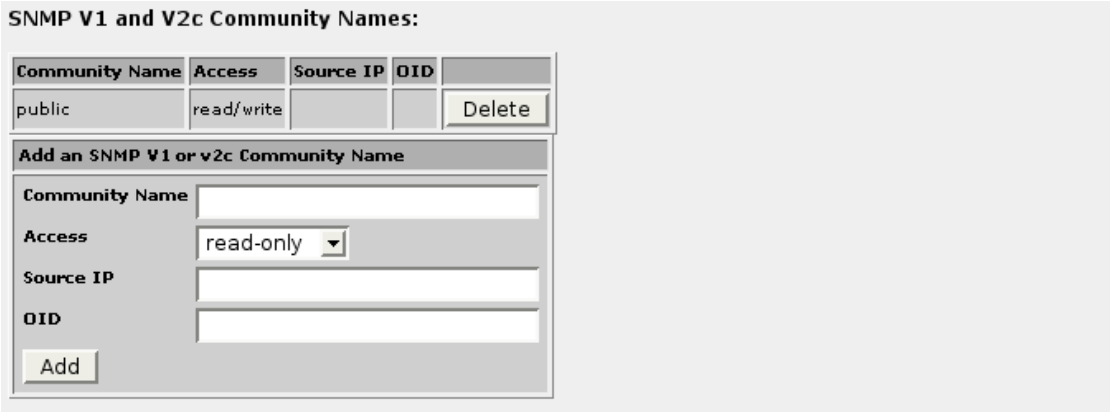

Figure 221: Access Control page, SNMP V1 and V2c

The first part of the Access control page allows the creation and deletion of SNMP V1 and V2c community names.

The **Community Name** field selects the name of the community. The **Access** field determines whether the community is read-only or read/write. The **Source IP** field may be used to specify an IP address or range (e.g. 10.0.0.0/24) from which access to this community name may be made. The **OID** field further restricts access to an Object Identifier (OID) tree at or below a specified OID.

Figure 222: Access Control page, SNMP V3

The second part of the Access control menu allows creation and deletion of V3 users.

The **User Name** field selects the name of the new user.

The **Access** field determines whether the community is read-only or read/write.

The **Minimum Security** field selects the level of security used by this user. It may be No Authentication (no authentication or encryption), Authentication Only (authentication by MD5 or SHA1 authentication methods, no encryption) or Authentication with Privacy (authentication by MD5 or SHA1, encryption by DES or AES ciphers).

The **OID** field further restricts access to an Object Identifier (OID) tree at or below a specified OID.

The **Authentication Protocol**, **Authentication Passphrase**, **Privacy Protocol** and **Privacy Passphrase** fields configure the protocols and passphrases used depending on the **Minimum Security** field. These settings are shared between agent and remote user.

Note that if authentication and privacy are both used, but only the authentication passphrase is provided, snmpd will use the authentication passphrase as the privacy passphrase.

Note also that if any notifications are enabled, a read-only user named **internal** will be automatically created to satisfy the requirements of the event MIB.

Trap Configuration

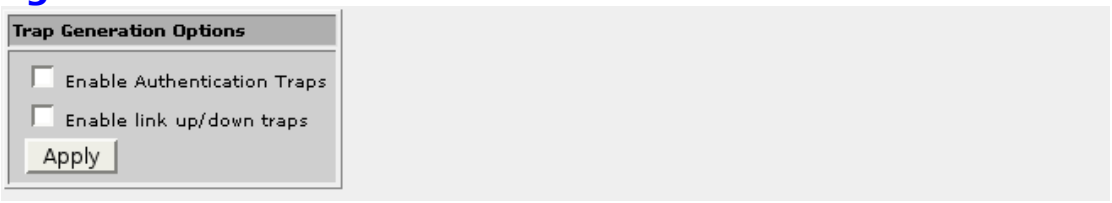

Figure 223: Trap Configuration page, Trap Options

The Trap Configuration page manages SNMP trap destinations. Under **Trap Generation Options**, you may enable the generation of notifications on authentication failures or IP interface link up/down events.

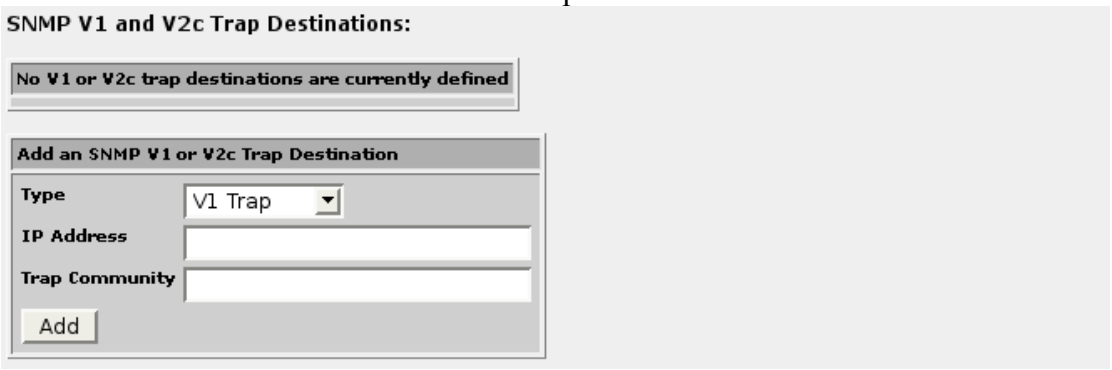

Figure 224: Trap Destinations V1 and V2c

The **SNMP V1 and V2c Trap Destinations** part of the menu allows the creation and deletion of trap destinations.

The **Type** field specifies the exchange used with this destination, either V1 trap, V2c trap or V2c inform.

The **IP address** and **Trap Community** fields specifies the receivers IP address and community name.

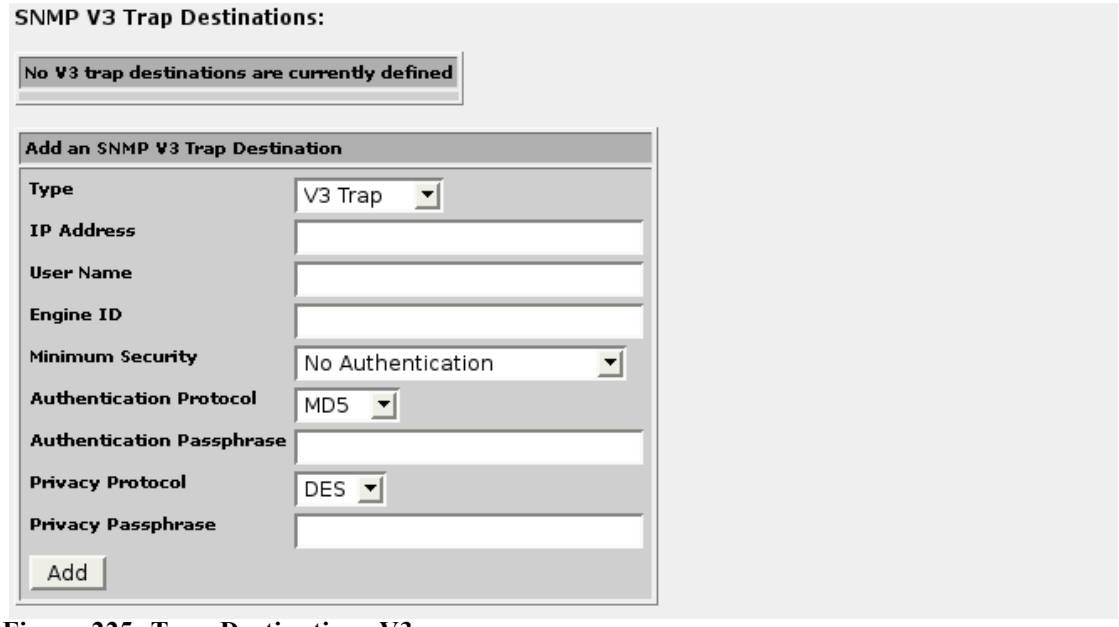

Figure 225: Trap Destinations V3

The **SNMP V3 Trap Destinations** part of the menu allows the creation and deletion of V3 trap destinations.

The **Type** field specifies the exchange used with this destination, either V3 trap or V3 inform.

The **IP address** and **Trap Community** fields specifies the receivers IP address and user name.

The **Engine ID** parameter is necessary for inform type notification destinations only, and must be configured by the trap receiver in order to receive these notifications.

The **Minimum Security**, **Authentication Protocol**, **Authentication Passphrase**, **Privacy Protocol** and **Privacy Passphrase** fields are as described above.

MIB Support

The RuggedRouter supports the following MIBs.

MIB Name	MIB Description
IF-MIB	The MIB module to describe generic objects for network interface sub-layers.
SNMPv2-MIB	The MIB module for SNMPv2 entities.
TCP-MIB	The MIB module for managing TCP implementations.
IP-MIB	The MIB module for managing IP and ICMP implementations.
UDP-MIB	The MIB module for managing UDP implementations.
SNMP-VIEW-BASED-ACM-MIB	View-based Access Control Model for SNMP.
SNMP-FRAMEWORK-MIB	The SNMP Management Architecture MIB.
SNMP-MPD-MIB	The MIB for Message Processing and Dispatching.
SNMP-USER-BASED-SM-MIB	The management information definitions for the SNMP User-based Security Model.

Radius Authentication

The Radius protocol described in RFC 2865 provides a means for carrying authentication, authorization, and configuration information between a client (the router) which desires to authenticate its links and a shared Authentication Server.

Transactions between the router and RADIUS server are authenticated through the use of a shared secret, which is never sent over the network. In addition, any user passwords are sent encrypted between the router and RADIUS server, to eliminate the possibility that someone snooping on an insecure network could determine a user's password.

Radius deals with categories of authentication, known as *services*. The router supports user logins via the **LOGIN** service, PPP connections via the **PPP** service and non-root Web management via the **WEBMIN** service. The WEBMIN service allows operator actions to be logged under their login name (as opposed to "root").

The router uses Radius to authenticate:

- Serial port, embedded modem and SSH console logins to the root account,
- SCP and SFTP (SSH file copies and file transfers) to the root account,
- Logins to the rrsetup configuration (rrsetup account),
- PPP Incoming connections on the embedded modem (specific user accounts),
- Web Management logins (root and radius user accounts).

Radius server redundancy is supported. Multiple Radius servers, usually operating from a common database, may be used to authenticate a new session. If the first configured Radius server does not respond, subsequent servers will be tried until a positive/negative acknowledgment is received or all servers have been tried.

Each server is configured with an associated timeout which limits the duration of the request to it. An authentication request could thus require up to the sum of the timeouts of all configured servers.

If no Radius servers are configured (or are able to authenticate the request), logins are authenticated from the system account stored on the router. The goal of Radius Authentication is usually to severely restrict the distribution of this password, limiting regular access to server based authentication.

Note: Users employing the WEBMIN service are the exception to this rule. Being entirely managed via radius, they cannot access web management if radius is down.

The user has the option of designating specific servers to authenticate either Logins, PPP or Webmin sessions or to have one server authenticate combinations of service or all services.

The radius server providing the WEBMIN service must also be configured to supply a "privilege-level" field which will be used in upcoming releases to provide operator levels of privilege. See the appendix on Radius Server Configuration for more information.

Helpful Hint

Some users set the rrsetup and root account passwords to difficult to guess strings that are unique to each router, then employ a common password for all routers in radius. The router specific strings are restricted to a very few personnel. A larger set of expert users are granted the rights to SSH login using the radius root account passwords. Yet another set of users are granted access via Webmin user accounts.

Radius authentication is logged to the authorization log (file auth.log). Details of each authentication including time of occurrence, source and result are included.

Radius Authentication Configuration

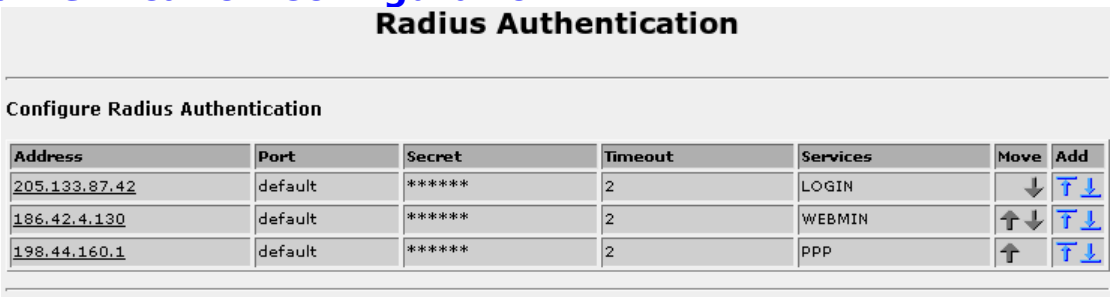

Figure 226: Radius Authentication Main Menu

Radius Authentication is configured from within the the **Maintenance** menu **Miscellaneous** sub-menu. This menu allows you to add, delete and Radius servers. Add a server by by clicking on the add-above or add-below arrows in the **Add** field. You may also edit a server by following its link under the **Address** field.

Reorder the servers by clicking on the arrows under the **Move** field.

Edit Radius Server Parameters

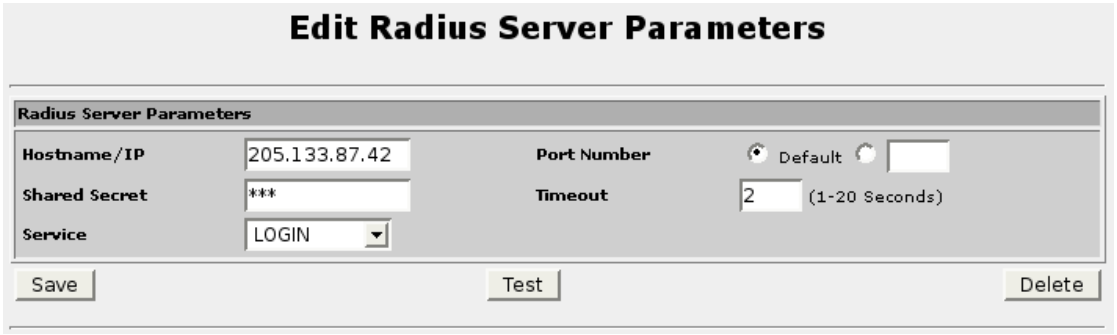

Figure 227: Radius Authentication Server Parameters

This menu configures, tests and deletes radius server entries.

The **Hostname/IP** field configures the server IP address.

The **Port Number** fields selects the default port number of 1812 or selects another specific port.

The **Shared Secret** field configures the unique password used by this server.

The time **Timeout** field selects the maximal time to wait before trying the next server.

The **Service** field configures whether the server authenticates LOGIN, WEBMIN, PPP LOGIN or any combination of these types.

Outgoing Mail

Figure 228: Radius Authentication Main Menu

Outgoing Mail is configured from within the the **Maintenance** menu **Miscellaneous** sub-menu. This menu controls where emails originated by the router are forwarded to.

The **Forward to Mail Hub** field specifies an IP address or domain name of a host that accept mail from the router.

The **Belongs to Domain** field specifies the email domain the router is part of. This information is written into the email header upon transmission.

The **Hostname** field specifies the hostname to be written into the email header upon transmission.

Helpful Hint
You can generate emails from scheduled commands and scripts with
“(echo "To: ops@myco"; echo -e "Subject: Hello!\n"; some-command) | sendmail -t”.

Chassis Parameters

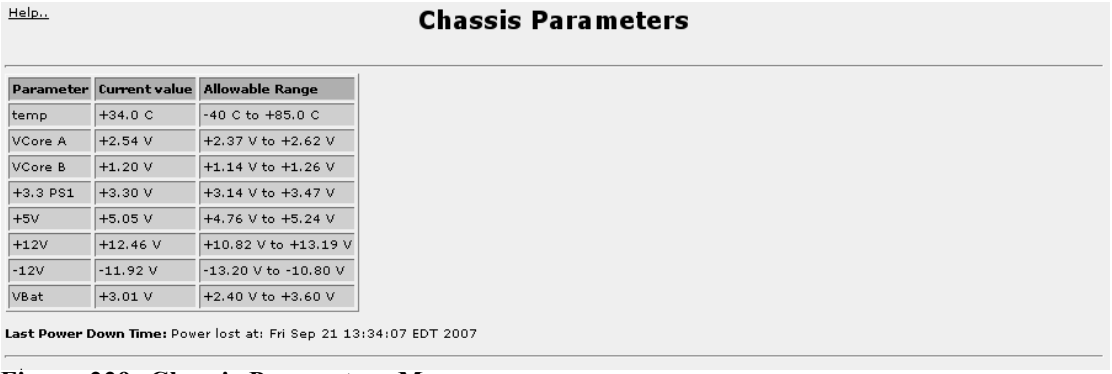

Figure 229: Chassis Parameters Menu

This menu displays the chassis temperature and, if hardware version 2, the voltage levels of chassis power supplies and a record of the last power down time. The system will highlight red any out-of-range value. The monitored values are described below:

Parameter	Description
temp	Motherboard temperature
VcoreA, VCoreB	Redundant 3.3V power supply voltages
+3.3 PS1, +3.3 PS2	Redundant 3.3V power supply voltages
+5V	5V power supply voltage
+12V	12V power supply voltage
VBat	Battery voltage

The last power down time reflects the time power was removed from the chassis as a result of a power failure, commanded reboot or an watchdog initiated reboot.

System alarms will be generated for out-of-range parameters and watchdog initiated reboots.

System Logs

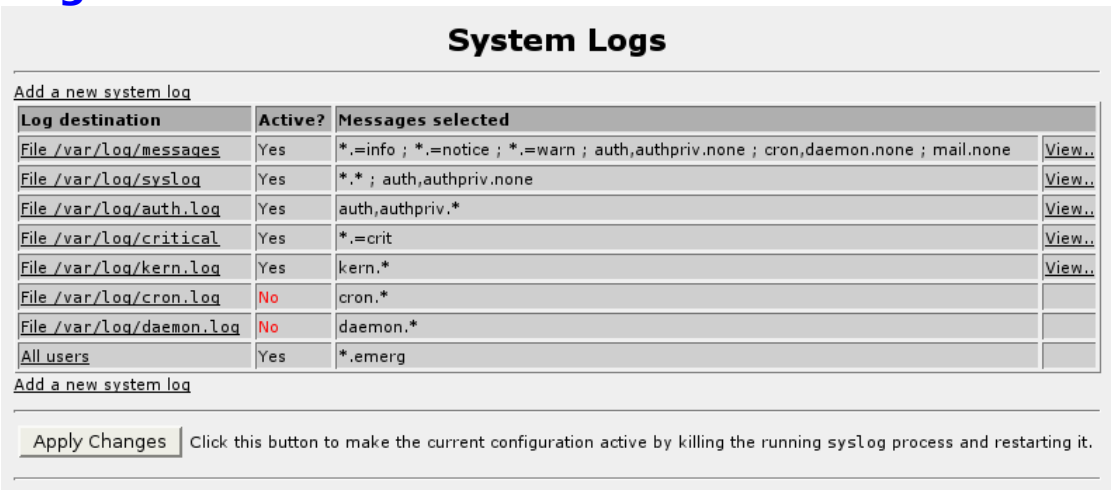

Figure 230: System Logs

System logs are records of activities that have occurred on the router, sorted into specific categories. System logs can be invaluable when debugging configuration changes. As such, most of your use of the logs will be simply in viewing them.

Syslog Factory Defaults

Although new logs can be created (and the type of information saved in existing logs changed) the factory defaults are as follows:

- messages – This log file catches a wide variety of generic information excluding authentication, cron and mail messages. This should be the first log you inspect when starting to debug a problem.
- syslog – This log file catches all information with the exception of authentications. Syslog contains all that messages contains, and more. Examine this log if you can not find relevant information in messages.
- auth.log – This log file catches authentication requests. View auth.log when you are trying to debug a problem in which a user is not able to sign on to a service (such as web management or ssh).
- critical – This log catches reports of critical failures. There should never be any messages in this log. Your RuggedCom support representative may ask you to inspect this file.
- kern.log – This log contains messages issued by the kernel (the most central part of the operating system). This log always displays messages issued at boot time, and should rarely be added to after that. Your RuggedCom support representative may ask you to inspect this file.
- cron.log (initially disabled) – This log file contains messages from the cron systems notifying of tasks started through cron. Your RuggedCom support representative may ask you to enable and inspect this log.
- daemon.log (initially disabled) – This log file contains messages from daemons (programs that run continuously in the background). Your RuggedCom support representative may ask you to enable and inspect this log.

Left unrestricted the logging system would consume all available “disk” space, causing the router to fail. The router limits the memory used by the logging system by storing logs in a volatile (i.e. lost after a reboot) file system which is limited in size. Such a system will lose logging information when a power failure occurs, too much logging is generated or as the result of a user commanded reboot.

The router deals with this problem by storing compressed versions of three key files (messages, auth.log, and critical) to the permanent disk. The log files are saved every 180 seconds and upon an orderly reboot. The log files are restored during the next boot. All other files but these are cleared.

Remote Logging

Remote logging (often referred to as remote syslogging) is the process of forwarding log entries to a remote host computer. Remote logging enables central collation of logs and preserves logs in the events of security incidents. Remote logging does not require any file storage on the router and as such does not suffer from loss of information around unplanned power failures. On the other hand, remote logging cannot record events that occur before network connectivity to the logging host is established.

Remote logging can replace disk logging or can augment it.

If you wish to replace disk logging for some information type, select the appropriate link under the **System Logs** sub-menu **Log Destination** column. Enter the URL of the logging host under the **Syslog server on**.

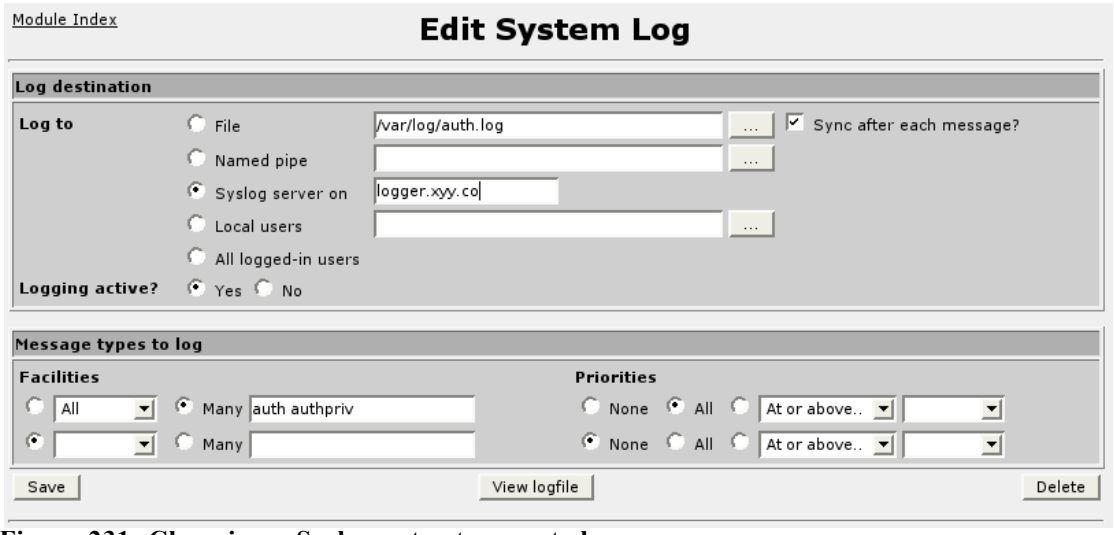

Figure 231: Changing a Syslog entry to remote log

If you wish to remote log in addition to disk log some log type, you must duplicate the log entry and the configure the logging host. Duplicate the entry by using the “Add a new system log” link on the **System Logs** sub-menu.

Finally, you may forward all information to the remote logger by creating a new system log entry and specifying “All” Facilities and all priorities, and checking the **Syslog server on** field with an appropriate address.

Upgrade System

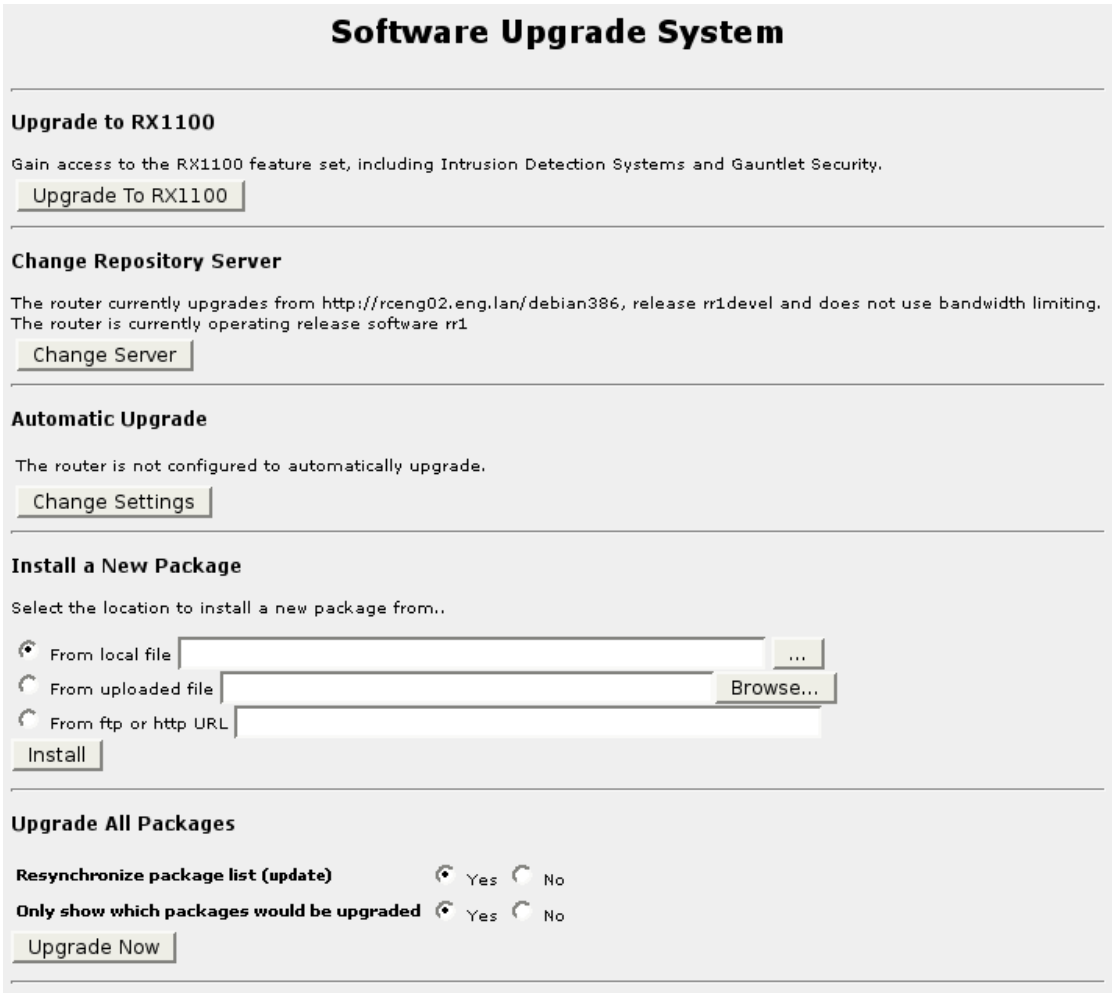

Figure 232: Software Upgrade System

The Software Upgrade system provides the following features:

- Upgrading from either HTTP or FTP servers,
- Upgrade traffic bandwidth limiting to prevent disruption to mission critical applications,
- Automatic daily upgrades from a central server at a scheduled time,
- Manually initiated upgrades from a central server,
- Manually initiated upgrades of new versions for testing purposes,
- Manually initiated installs of new packages for testing purposes.

RuggedRouter Software Fundamentals

You may be required to upgrade the router in order to take advantage of new features, security improvement and bug repairs.

Your RuggedRouter software is provided in releases of the form rrX.Y.Z. The platform release number X changes when new hardware platforms are released. The major release number Y is increased when important new features are added. This is called a “Major” release. The minor release number Z is increased when minor functionality is added or bug repairs are made. This is called a “Minor” release.

The actual software of the RuggedRouter is composed of a number of “packages”. Each package contains all of the files necessary to implement a set of related commands or features, such as a firewall or ssh client. A router upgrade involves replacing some of these packages with newer versions and with adding new packages. The upgrade system handles all this for you.

When A Software Upgrade Requires A Reboot

Software release upgrades that involve changing to a new linux kernel require a reboot. Releases that force a reboot are always “Major” releases (but note that some major releases may not require a reboot). Minor releases will never require a reboot. The upgrades release notes will state whether a reboot will occur.

You are warned before starting an upgrade when a reboot is required. The only exception is the unattended automatic upgrade.

Automatic Upgrade

It can be programmed to check a server on your network at a specific time each day, upgrading to the newest release. RuggedCom understands that some administrators may wish to pre-test package upgrades on specific machines before performing a network wide upgrade. It is also possible to manually control the upgrade process on a per-machine per-package basis.

The upgrade system allows you to restrict the maximum amount of bandwidth consumed by an upgrade. Most upgrades will involve relatively modest amounts of data transfer, especially over an Ethernet class network. But when the router is accessed over a low speed WAN link, even a small upgrade can temporarily consume 100% of the links bandwidth. This can disrupt mission critical applications. The bandwidth limiting feature limits only the upgrading process, leaving regular traffic unregulated.

Upgrade to RX1100

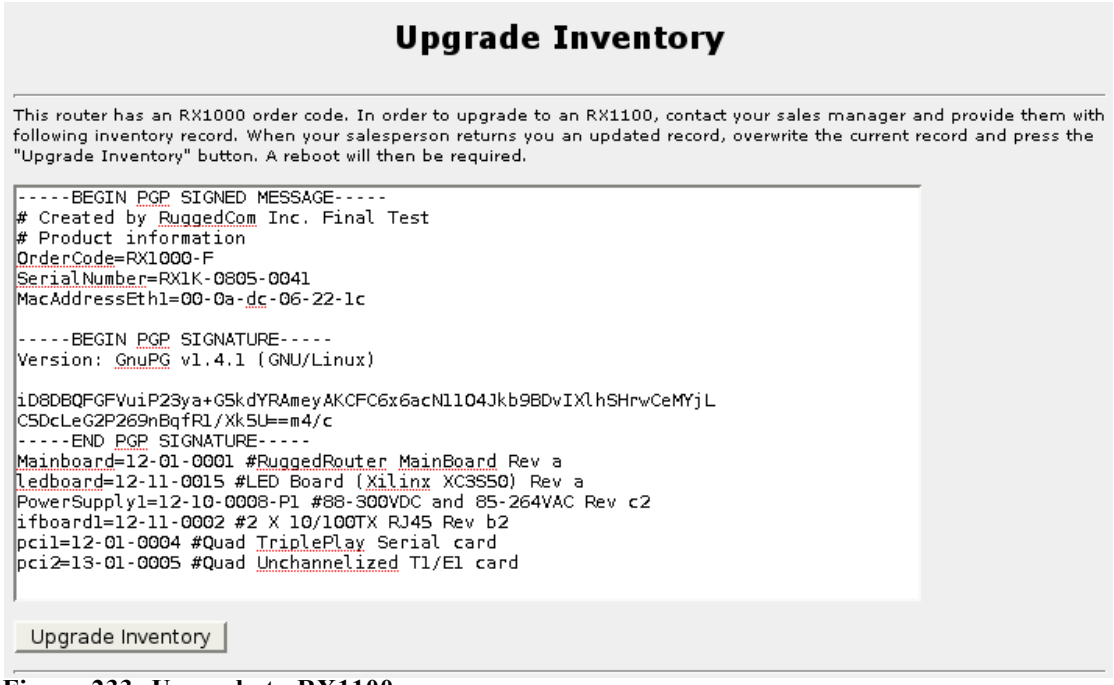

Figure 233: Upgrade to RX1100

This menu allows you to upgrade your router. The display usefully provides a description of the current hardware in the router inventory.

Change Repository Server

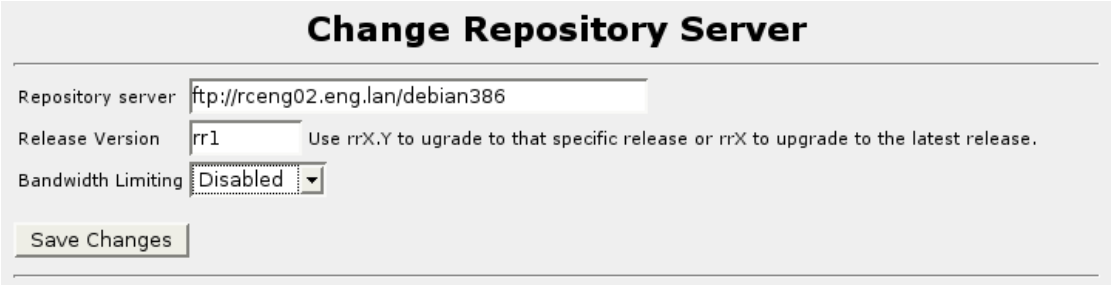

Figure 234: Change Repository Server

This menu defines the server used to upgrade software. The **Repository server** field accepts a URL containing the domain name or IP address of an http or ftp server along with the directory on the server containing the upgrades.

The **release version** field accepts a software release string, such as “rr1” or “rr1.7” or rr1.7.2.

If you configure this field with only a major release number such as “rr1”, the router will always pick the latest release at the server. As an example, if the router is running with release rr1.7 and release rr1.7.2 becomes available, the latter will be used.

If you configure this field with a major/minor/patch release number such as “rr1.7.2”, the router will only upgrade from that release.

The **Bandwidth Limiting** selector allows you to select the bandwidth available for upgrading software.

Automatic Upgrading

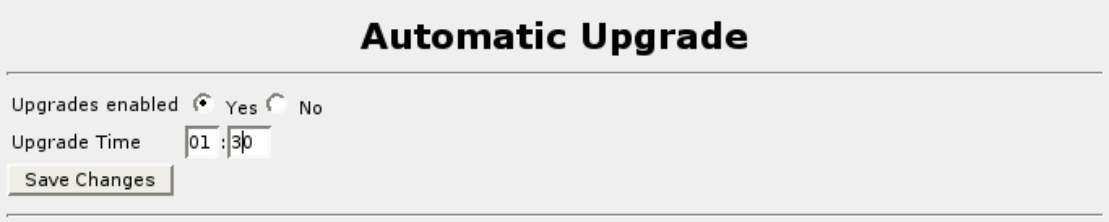

Figure 235: Automatic Upgrade

Check the **Upgrades enabled** field to activate daily upgrades.

Use the **Upgrade Time** fields to select the time to upgrade. Selecting different times on each router can be used to even out traffic flows in the network.

Upgrading All Packages

Figure 236: Upgrading All Packages

The Upgrading All Packages feature works by obtaining a list of the latest packages and then either showing what needs to be upgraded or actually doing an upgrade.

The R**esynchronize package list** field selects whether to obtain the list. You only need to obtain the list once, so checking **No** can save you some time if your first pass was "Only show". This is especially true if the network link is a low-speed WAN link.

The **Only show which packages would be upgraded** field controls whether to show what needs to be upgraded or actually do an upgrade.

Note: *Webmin manages the upgrade of other packages. When Webmin must upgrade itself, the process requires an extra step. You will be requested to start a Webmin only upgrade. Webmin will start another program to manage the upgrade and will self-terminate. Webmin will automatically restart after the upgrade completes, after which time you may log back in.*

Installing A New Package

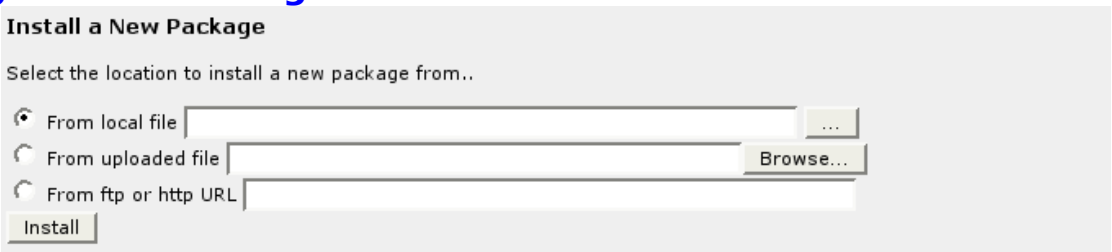

Figure 237: Installing A New Package

The Install A New Packages feature uploads and installs packages to the router.

Select the **From local file** option if you have already moved the package to the router through http, ftp or scp. You may either enter the full path from the root directory to the package or use the file selector (...) to identify the package.

Select the **From uploaded file** option if you have the file locally on your workstation. You may either enter the location of the file on your local file system browse selector (Browse...) to identify the package.

Select the **From ftp or http URL** if you know the network address of the package.

Complete the installation by selecting the install button.

Pre-upgrade/Post-upgrade scripts

The pre-upgrade and post-upgrade script feature allows you to execute a user defined script before and after a software system upgrade.

The scripts run only when there are packages to install, they will not run when "showing" packages that could be installed.

During a real upgrade, the router will try and download the scripts from the same location as is configured by the **Change Repository Server** page. The router will attempt to download a file named "pre-upgrade" and execute it before the upgrade starts. After the upgrade completes (including webmin) the router will attempt to download a file named "post-upgrade" and execute it.

The scripts start with #!/bin/bash or #!/usr/bin/perl and be designed to produce consistent results in the event they are subsequent run a second time. It is possible that the upgrade can be interrupted after the preupgrade script runs, and re-started at a later date.

The result of running the pre-upgrade script is included in the upgrade output. If ran through the automatic upgrade, the scripts output can be viewed by through the "View Log File of Last Upgrade" button on the Software Upgrade System page.

Example of a post-upgrade script: The following post-upgrade script will send an email notification when upgrade completes (assuming ssmtp is configured properly).

```
#!/bin/bash
echo "Subject: Software upgrade for Release rr1.9.0 on `hostname` completed" > /tmp/mail
echo "To: controlcenter@ruggedcom.com" >> /tmp/mail
echo "Software upgrade for Release rr1.9.0 on `hostname` completed at `date`" >> /tmp/mail
echo >> /tmp/mail
cat /tmp/mail | ssmtp controlcenter@ruggedcom.com
rm -f /tmp/mail
```

Uploading And Downloading Files

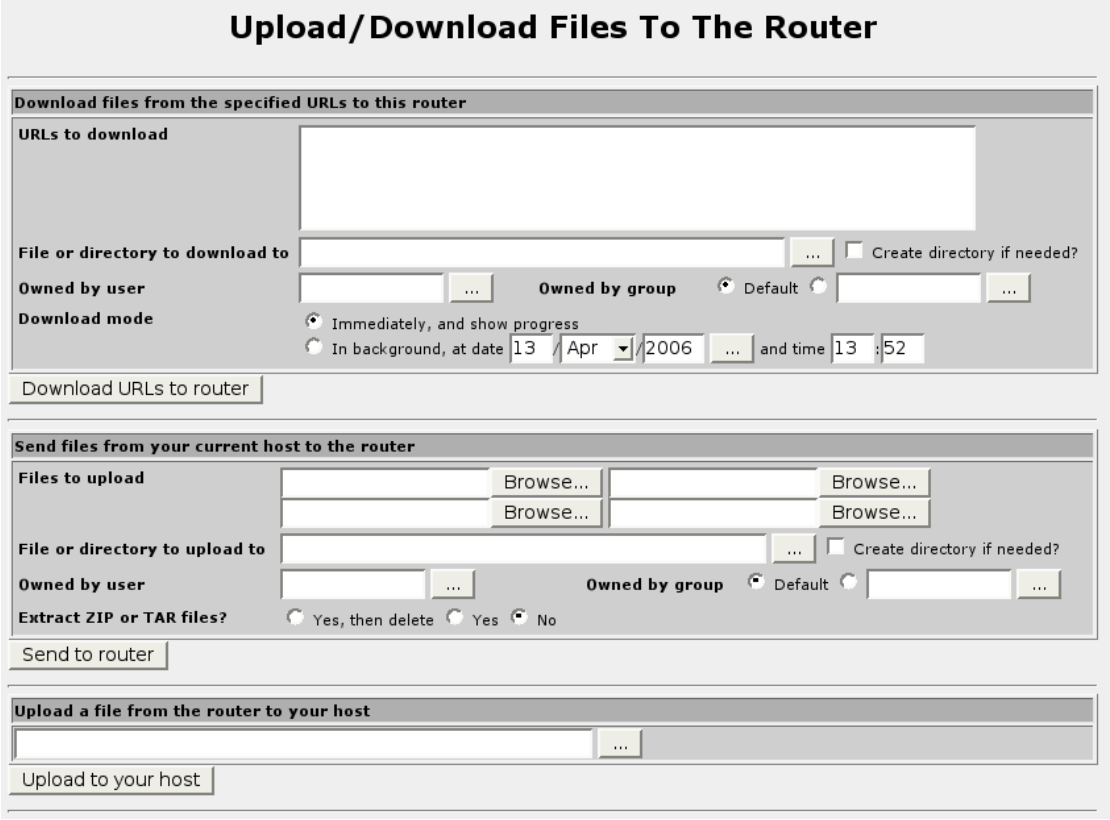

Figure 238: Upload/Download menu

The Upload/Download Files menu provides a means to transfer files to and from the router.

The **Download files from the specified URLs to this router** part of the menu allows you to have the router download files from **ftp** and **http** servers. You need to specify (at least) the file URL and the directory to download it to. You may also decide to create directories cited in the download path at download time, set the user/group ownership of the file and postpone the download to a specific time.

The **Send files from your current host to the router** part of the menu allows you to send files from your host machine directly to the router. You need to specify (at least) one file to send and the directory to upload it to. Clicking on a browse button will open a file search dialog box. Select the file to upload to the router and close the dialog box. Click upon the **Send to router** button to start the transfer. You may also decide to create directories cited in the upload path at upload time, set the user/group ownership of the file and extract tar or zip files.

The **Upload a file from the router to your host** part of the menu allows you to send files from the router a specified your host machine. You need to specify the file to send. You may specify the files path directly or click on the browse button to open a file search dialog box. Select the file to upload and close the dialog box. Then click the **Upload to your host** button.

Chapter 27 – Security Considerations

Introduction

This chapter describes actions to take to secure the RuggedRouter.

Security Actions

1. Change the root and rrsetup passwords from the rrsetup shell, before attaching the router to the network.
2. If Radius authentication is being employed, configure authentication servers.
3. Restrict the IP addresses which Web management will accept connections from. See the **Webmin** menu, **IP Access Control** sub-menu. Restrict the Ethernet ports which Web management will accept connections from. See the **Webmin** menu, **Ports and Addresses** sub-menu.
4. Review the IP networking settings provided in the **Network Configuration** menu, **Core Settings** sub-menu. You may wish to tighten some settings, especially Ignore All ICMP ECHO requests.
5. Restrict the users that the SSH server will allow to connect. See the **SSH Server** menu, **Access Control** sub-menu.
6. If the router is an RX1100 and you wish to use the Snort Intrusion Detection System, activate and configure it.
7. If the router is an RX1100 and you wish to use the Gauntlet security appliance, activate and configure it.
8. If SNMP will be used, limit the IP addresses which can connect and change the community names. Configure SNMP to raise a trap upon authentication failures.
9. Only enable the services you need and expect to use.
10. The RuggedRouter comes with the following login banner. Replace the contents of the file /etc/issue and /etc/issue.net in order to change it.

```
WARNING: You are attempting to access a private computer system. Access to this
system is restricted to authorized persons only. This system may not be used for
any purpose that is unlawful or deemed inappropriate.  Access and use of this
system is electronically monitored and, by entering this system, you are giving
your consent to be electronically monitored. We reserve the right to seek all
remedies for unauthorized use, including prosecution.
```

11. If using a firewall, configure and start the firewall before attaching the router to the public network. Configure the firewall to accept connections from a specific domain.
12. Configure remote system logging to forward all logs to a central location.

This page intentionally blank

Appendix A – Setting Up A Repository

The RuggedCom software upgrade mechanism requires a repository of software to available. The following instructions detail:

- Requirements for a repository server,
- Initial set up of a repository,
- Upgrading the repository to the latest release,
- Maintain separate releases streams for different groups of routers,
- Setting up one router to test new releases
- Configuring the network routers.

Repository Server Requirements

In order to establish a repository you will need a host that is accessible to the routers that will be upgraded. This host must be able to act as a web server or ftp server. The host must also be able to access the RuggedCom web site in order to download new releases of software from RuggedCom.

The server requirements are fairly modest. The principal requirements are for disk space, bandwidth and the ability to serve an adequate number of http sessions.

Each software release will require approximately 50 Mb of disk space. Note that this figure includes an entire software image, most upgrades will involve the transfer of only a small fraction of this amount. A large number of such releases could easily be stored on a system of only modest capabilities. In practice, only one or two releases are usually all that need be kept.

The bandwidth requirements are determined by the many factors including the number of routers, size of upgrade, when the routers upgrade, bandwidth limiting at each router and network bandwidth capability. Most web servers can serve files to the limit of the network interface bandwidth, so even a modest (e.g. 486 class machine) would prove acceptable.

The server should be able to accept at least as many http or ftp connections as there are upgradable routers in the network. In practice you will configure the routers to have staggered upgrade times in order to minimize the impact of upgrading on the network. A large upgrade (or a low bandwidth limiting value at each router) may cause all the routers to be upgrading at any one time.

Initial Repository Setup

You must create a directory on the web server to hold the releases for the router. The directory can have any name, such as “ruggedrouter”.

Some administrators like to designate one router to test the impact of new software. This will require a directory, such as “ruggedroutertest” to be created.

These directory names will be used in examples in the remainder of this section.

Ensure that the web server publishes these directories.

Upgrading The Repository

RuggedRouter releases are obtained from the RuggedCom web site as ZIP files. Download the ZIP file to your regular and/or test release directories and unzip them. You may delete the original ZIP file if desired.

The ZIP file name will be in the form rrX.Y.zip. The major release number X is changed when major new functionality (often hardware related) is offered. The minor release number Y is increased when minor functionality is added or bug repairs are made. The first RuggedRouter upgrade release is rr1.1.zip.

The zip file will extract to a directory that has the same name as the major release, e.g. “rr1”. As subsequent release are made, they will also be extracted into this directory.

Setting Up The Routers

The name of the release directory, and the major and minor release names from the zip file tells you how to set up the routers.

Suppose you have just unzipped rr1.2.zip into “ruggedroutertest” on a server available to the network at server.xyz.net. The major release is rr1 and the minor release is 2. You have chosen this directory because you want to test the release on a specific machine before propagating it to the network.

Login to the test router and visit the **Maintenance** menu, **Upgrade Software**, **Change Repository Server** sub-menu. Change the **Repository server** field to “http://server.xyz.net/ruggedroutertest” and the **Release Version** field to “rr1”. You can proceed to upgrade the router manually or wait for the next nightly upgrade to take place.

After you are satisfied that the upgrade was successful you can proceed to unzip the rr1.2.zip file into your “ruggedrouter” directory (or copy the rr1/dists/rr1.2 and rr1/dists/current directories into or the “ruggedrouter” directory).

Ensure that the remainder of the routers to be upgraded have a **Repository server** field to “http://server.xyz.net/ruggedrouter” and the **Release Version** field to “rr1”. They can now be upgraded.

An Alternate Approach

You can eliminate the need for separate release and test directories by making your routers upgrade to a specific major and minor releases.

In this approach you will always extract releases to the same directory, e.g. “ruggedrouter”.

All routers will be configured with a **Repository server** field set to “http://server.xyz.net/ruggedrouter” and the **Release Version** field initially set to “rr1.1”. When you need to upgrade to rr1.2 you will visit the routers and update the **Release Version** field.

This method is simpler, but has the disadvantage that you need to visit each of the routers. This can become unwieldy when there are many routers to manage.

Upgrading Considerations

The RuggedRouter offers you the ability to perform automatic daily upgrades, specify the download time and limit the download bandwidth. These tools automate the upgrade process and minimize the impact of upgrading on the network.

When automatic daily upgrades are used, you may wish to stagger the upgrade time of the routers. If your network has a natural “ebb flow” period of traffic activity, schedule the upgrades during this time. As an example if you have 20 routers to upgrade and they must be upgraded over an eight hour period, configure each router to start its upgrade 20 minutes after the previous router.

Be careful with limiting download bandwidth in the router. Typical upgrades will involve less than 5 MBytes of traffic. If bandwidth limiting is employed and set to 8 Kbps the upgrade will require upwards of 1.5 hours to complete.

Administrators should also be wary of routers which concentrate locally connected routers as the upgrade bandwidth consumed on the network link could reach the sum of all bandwidth limiting settings.

Routers using Frame Relay with CIR under-subscription may also encounter lengthier downloads because of retransmission.

Appendix B – Downgrading Router Software

RuggedCom recognizes that customers may need to downgrade router software:

- Routers being added to the network have more recent version than that standardized for the network.
- Network staff may wish to regain confidence in the software of an exposed router by downgrading it to its current version, essentially reloading its software.
- Network staff may wish to explore how features operated on a previous release.

The release process involves the following steps:

1. The downgrade image file is downloaded from RuggedCom to a web server.
2. The router to downgrade is attached via one of its Ethernet ports to the web server (either directly or via a network), configured and tested.
3. The router can also load a specific configuration archive on to the target router. This file is also loaded on the web server.
4. The router is rebooted and is forced to enter its boot selection menu by pressing the down arrow key of an attached terminal continuously. The router will offer a menu that will provide the option "Software Downgrade Utility". Select this option and press enter.
5. The router will prompt for the required information.
6. The router will downgrade the software and reboot. This will require five to ten minutes to complete.

Note: The router must not lose power or be interrupted during the downgrade process. *The process involves a complete rewrite of the operating image. Interruption will require that the router be returned to the factory to have the software restored. RuggedCom suggests minimization of problems by using a standalone PC as the web server and by powering both the PC and the router from an uninterruptible power source (UPS).*

The following information will be requested during the download:

- The Ethernet port to use, whether to use DHCP or assign a static IP address and a network mask.
- A gateway IP address (if one is needed).
- The URL of the image file on the web server.
- The URL of the configuration archive file on the web server (if one is used).

Appendix C – Installing Apache Web Server On Windows

A number of customers have asked for advice and instructions on setting up a web server on Windows. RuggedCom recommends the Apache web server, because it is secure, robust, easy to install and configure as well as being able to be installed on a wide variety of Windows platforms.

Begin by identifying a host computer and its physical and logical location on the network. The **Repository Server Requirements** of the appendix "Setting Up A Repository" provide some guidance on host requirements. The Apache installation process will prompt you for an IP address and domain name with which to serve the web pages. Later in the install, you will also need to provide the directory where the RuggedRouter releases will be kept. Ensure that a web servers is not already installed.

Obtain Apache by visiting the web page of www.apache.org. Visit the "HTTP Server" portion of the web site and click on the "Downloads" page. Identify the latest version of Apache and find its Win32 version, usually under "httpd/binaries/win32/". You should be able to find a Microsoft System Installer Version (e.g. apache_2.0.55-win32-x86-no_ssl.msi), as well as platform specific notes. Download and install this version.

Verify the web server by opening a web browser on another host on the network and entering the URL http:// followed by the IP address Apache was installed with. Note that you may also verify Apache from a browser on the web server itself by browsing http://localhost. If properly set-up, the Apache default web page will be shown.

Figure 239: Apache Default Web Page

Apache serves the web pages contained in the directory known as the "DocumentRoot". You must change the document root by, from the desktop, clicking Start -> All programs -> Apache HTTP Server -> Configure Apache Server -> Edit the Apache httpd.conf file. Search the file for the DocumentRoot variable and change it to the directory where your RuggedRouter release are kept. Restart Apache by clicking Start -> All programs -> Apache HTTP Server -> Control Apache Web Server -> Restart.

Return to the web browser used earlier to verify Apache and refresh the screen. It should now reflect the contents of your RuggedRouter release directory. You should now be able to perform an upgrade from a router.

Appendix D – Installing IIS Web Server On Windows

A number of customers have asked for advice and instructions on setting up an IIS web server on Windows.

Begin by identifying a host computer that has IIS and its physical and logical location on the network. The **Repository Server Requirements** of the appendix "Setting Up A Repository" provide some guidance on host requirements.

Start to install IIS by clicking on **Start menu, Control Panel, Add or Remove Programs, Add/remove Windows Components.** In the resultant menu check the **Internet Information Services(IIS)** box and select next.

Figure 240: Installing IIS

Apache serves the web pages contained in the directory known as th

Download the desired release (e.g. rr1.9.0.zip) from the RuggedCom website. Create the directory ruggedcom under the IIS root directory **C:\Inetpub\wwwroot**. Unzip the rr1.9.0.zip file within **C:\Inetpub\wwwroot\ruggedcom.**

Start to enable IIS **by c**licking on **Start menu, Control Panel, Administrative Tools, Internet Information Services**. Right click on **Internet Information Services, Connect** and enter the host computer's IP address, e.g. 192.168.0.1.

Verify the IIS web server by opening a web browser on another host on the network and entering the URL http:// followed by the IP address IIS was installed with, followed by /ruggedcom, e.g. http://192.168.0.1/ruggedrouter.

Visit the router you wish to upgrade and visit the **Maintenance** menu, **Upgrade System** sub-menu. Click on the **Change Server** button and set the Repository Server field (e.g. http:// 192.168.0.1/ruggedcom). Set the Release Version field to rr1. Save the configuration and return to the **Maintenance** menu. Set the **Only show which packages would be upgraded** radio button to **No** and click on the **Upgrade Now** button to start the upgrade.

Appendix E – Radius Server Configuration

This section describes how to configure popular radius servers to supply a Vendor-Specific field, “privilege-level”, which is used by Webmin to assign assign specific capabilities to Webmin users on a per user basis. Currently, the only privilege-level is that of “root”, but RuggedCom will be introducing additional levels in upcoming releases.

FreeRadius

The following steps to add Vendor-Specific attributes to the freeradius radius server.

1. Locate your dictionary file (usually in /usr/share/freeradius/).
2. In your dictionary directory, open the file “dictionary” add the line “$INCLUDE dictionary.ruggedcom” to the end of it
3. Create a file “dictionary.ruggedcom” under the dictionary directory containing:

```
# -*- text -*-
#
#   The RuggedCom Vendor-Specific dictionary.
#
# Version:  $Id: dictionary.RuggedCom,v 1.3.4.1 2005/11/30 22:17:24 aland Exp $
#
#   For a complete list of Private Enterprise Codes, see:
#
#   http://www.isi.edu/in-notes/iana/assignments/enterprise-numbers
#
VENDOR      RuggedCom           15004

BEGIN-VENDOR    RuggedCom

ATTRIBUTE   RuggedCom-Privilege-level           2   string

END-VENDOR RuggedCom
```

4. Users are assigned by adding lines to the file /etc/freeradius/user. Note that currently, the only privilege-level is that of “root”. For example to assign a user “john” with a password of “test”, add the following line:

```
john Auth-Type := Local, User-Password == "test"
    4.      RuggedCom-Privilege-level = "root"
```

5. Restart your freeradius server.

Windows Internet Authentication Service

The following steps to configure your IAS server.

1. Create groups used for different privilege level, for example, if the privilege level is root, you can create a group called Radius_RuggedRouter_root. Add the users having this privilege level to this group.
2. Use the New Remote Access Policy Wizard to create a custom policy with the following settings:

   Conditions:

   NAS-Identifier matches with webmin

   Windows-Group matches with the group the user belongs to

Permission: Grant remote access permission

3. Double click the policy name you created, In the popup window, click Edit Profile... button.

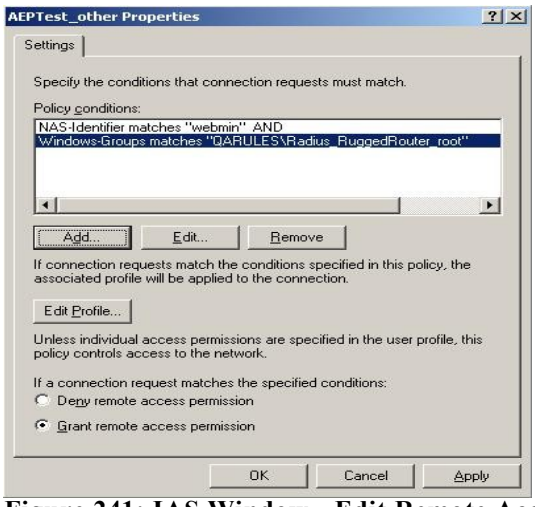

Figure 241: IAS Window - Edit Remote Access Policy

4. In Edit Profile window, Click Add... button

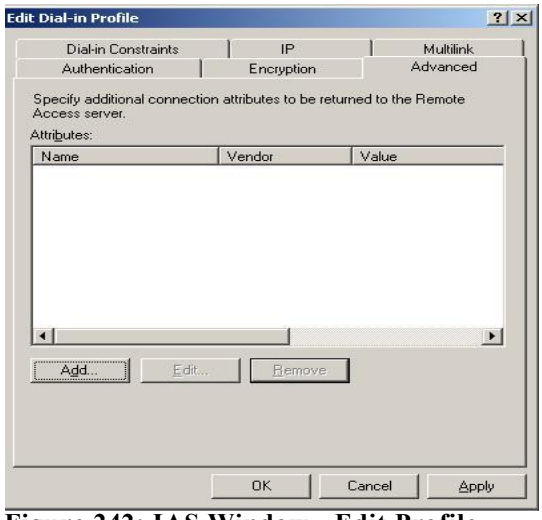

Figure 242: IAS Window - Edit Profile

5. In Add Attribute window, select Vendor-Specific line, and click Add button.

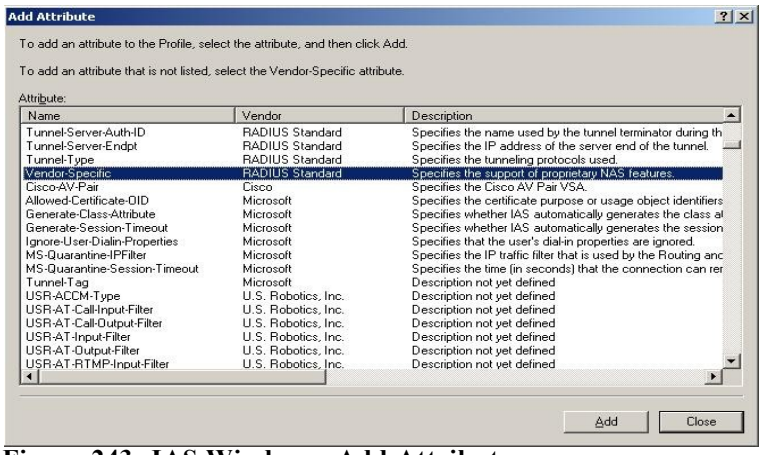

Figure 243: IAS Window – Add Attribute

6. In the Multivalued Attribute Information window, click the Add button

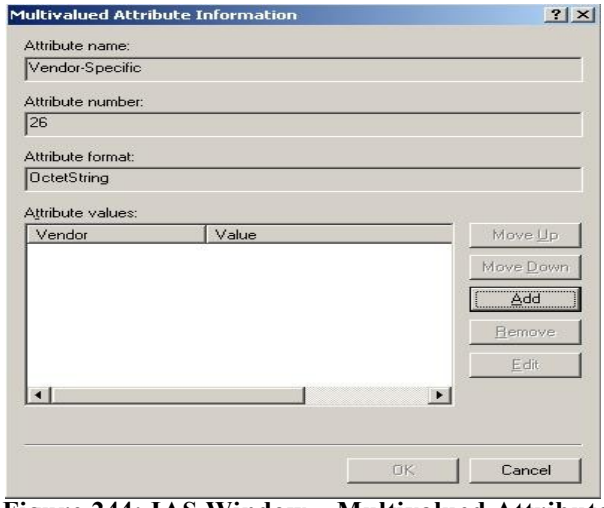

Figure 244: IAS Window – Multivalued Attribute Information

7. In the Vendor-Specific Attribute Information window, select radio button Enter Vendor Code, and input 15004 to the editbox. Select the radio button Yes, It conforms and click the button Configure Attribute...

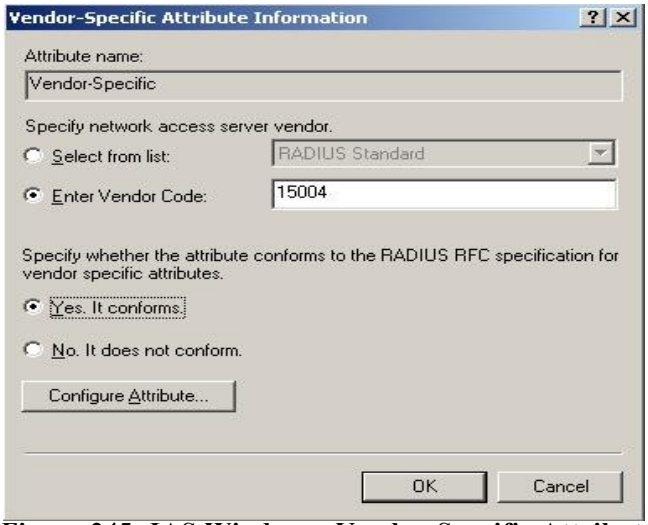

Figure 245: IAS Window – Vendor-Specific Attribute Information

8. In the Configure VSA (RFC compliant) window, in the vendor-assigned attribute number editbox, input 2; in the Attribute format listbox, select String, in the Attribute value editbox, input the desired privilege level (in the above case, it is operator, in your case, currently you should input root).

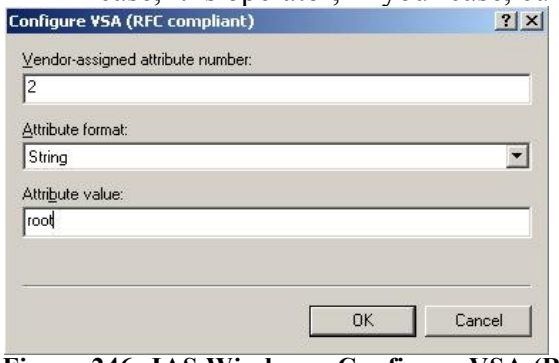

Figure 246: IAS Window – Configure VSA (RFC compliant)

Index